



OpenShift Container Platform 4.16

Installing

Installing and configuring OpenShift Container Platform clusters

OpenShift Container Platform 4.16 Installing

Installing and configuring OpenShift Container Platform clusters

Legal Notice

Copyright © 2024 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, the Red Hat logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux[®] is the registered trademark of Linus Torvalds in the United States and other countries.

Java[®] is a registered trademark of Oracle and/or its affiliates.

XFS[®] is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL[®] is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js[®] is an official trademark of Joyent. Red Hat is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack[®] Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

This document provides information about installing OpenShift Container Platform and details about some configuration processes.

Table of Contents

CHAPTER 1. OPENSIFT CONTAINER PLATFORM INSTALLATION OVERVIEW	70
1.1. ABOUT OPENSIFT CONTAINER PLATFORM INSTALLATION	70
1.1.1. About the installation program	70
1.1.2. About Red Hat Enterprise Linux CoreOS (RHCOS)	71
1.1.3. Glossary of common terms for OpenShift Container Platform installing	71
1.1.4. Installation process	73
The installation process with the Assisted Installer	74
The installation process with Agent-based infrastructure	74
The installation process with installer-provisioned infrastructure	74
The installation process with user-provisioned infrastructure	75
Installation process details	75
1.1.5. Verifying node state after installation	77
Installation scope	78
1.1.6. OpenShift Local overview	78
1.2. SUPPORTED PLATFORMS FOR OPENSIFT CONTAINER PLATFORM CLUSTERS	79
CHAPTER 2. SELECTING A CLUSTER INSTALLATION METHOD AND PREPARING IT FOR USERS	82
2.1. SELECTING A CLUSTER INSTALLATION TYPE	82
2.1.1. Do you want to install and manage an OpenShift Container Platform cluster yourself?	82
2.1.2. Have you used OpenShift Container Platform 3 and want to use OpenShift Container Platform 4?	83
2.1.3. Do you want to use existing components in your cluster?	83
2.1.4. Do you need extra security for your cluster?	84
2.2. PREPARING YOUR CLUSTER FOR USERS AFTER INSTALLATION	84
2.3. PREPARING YOUR CLUSTER FOR WORKLOADS	84
2.4. SUPPORTED INSTALLATION METHODS FOR DIFFERENT PLATFORMS	85
CHAPTER 3. CLUSTER CAPABILITIES	90
3.1. SELECTING CLUSTER CAPABILITIES	90
3.2. OPTIONAL CLUSTER CAPABILITIES IN OPENSIFT CONTAINER PLATFORM 4.16	92
3.2.1. Bare-metal capability	92
Purpose	92
3.2.2. Build capability	93
Purpose	93
3.2.3. Cloud controller manager capability	93
Purpose	93
3.2.4. Cloud credential capability	94
Purpose	94
3.2.5. Cluster storage capability	95
Purpose	95
Notes	95
3.2.6. Console capability	95
Purpose	95
3.2.7. CSI snapshot controller capability	95
Purpose	95
3.2.8. DeploymentConfig capability	95
Purpose	96
3.2.9. Ingress Capability	96
Purpose	96
Project	96
CRDs	96
Configuration objects	96

Notes	96
3.2.10. Insights capability	97
Purpose	97
Notes	97
3.2.11. Machine API capability	97
Purpose	97
3.2.12. Marketplace capability	97
Purpose	97
3.2.13. Operator Lifecycle Manager capability	98
Purpose	98
3.2.14. Node Tuning capability	98
Purpose	98
3.2.15. OpenShift samples capability	99
Purpose	99
3.2.16. Cluster Image Registry capability	99
Purpose	99
Project	100
3.3. ADDITIONAL RESOURCES	100
CHAPTER 4. DISCONNECTED INSTALLATION MIRRORING	101
4.1. ABOUT DISCONNECTED INSTALLATION MIRRORING	101
4.1.1. Creating a mirror registry	101
4.1.2. Mirroring images for a disconnected installation	101
4.2. CREATING A MIRROR REGISTRY WITH MIRROR REGISTRY FOR RED HAT OPENSIFT	101
4.2.1. Prerequisites	101
4.2.2. Mirror registry for Red Hat OpenShift introduction	102
4.2.2.1. Mirror registry for Red Hat OpenShift limitations	102
4.2.3. Mirroring on a local host with mirror registry for Red Hat OpenShift	103
4.2.4. Updating mirror registry for Red Hat OpenShift from a local host	104
4.2.5. Mirroring on a remote host with mirror registry for Red Hat OpenShift	105
4.2.6. Updating mirror registry for Red Hat OpenShift from a remote host	106
4.2.7. Replacing mirror registry for Red Hat OpenShift SSL/TLS certificates	106
4.2.8. Uninstalling the mirror registry for Red Hat OpenShift	107
4.2.9. Mirror registry for Red Hat OpenShift flags	108
4.2.10. Mirror registry for Red Hat OpenShift release notes	109
4.2.10.1. Mirror registry for Red Hat OpenShift 1.3.11	109
4.2.10.2. Mirror registry for Red Hat OpenShift 1.3.10	109
4.2.10.3. Mirror registry for Red Hat OpenShift 1.3.9	110
4.2.10.4. Mirror registry for Red Hat OpenShift 1.3.8	110
4.2.10.5. Mirror registry for Red Hat OpenShift 1.3.7	110
4.2.10.6. Mirror registry for Red Hat OpenShift 1.3.6	110
4.2.10.7. Mirror registry for Red Hat OpenShift 1.3.5	110
4.2.10.8. Mirror registry for Red Hat OpenShift 1.3.4	111
4.2.10.9. Mirror registry for Red Hat OpenShift 1.3.3	111
4.2.10.10. Mirror registry for Red Hat OpenShift 1.3.2	111
4.2.10.11. Mirror registry for Red Hat OpenShift 1.3.1	111
4.2.10.12. Mirror registry for Red Hat OpenShift 1.3.0	111
4.2.10.12.1. New features	111
4.2.10.12.2. Bug fixes	112
4.2.10.13. Mirror registry for Red Hat OpenShift 1.2.9	112
4.2.10.14. Mirror registry for Red Hat OpenShift 1.2.8	112
4.2.10.15. Mirror registry for Red Hat OpenShift 1.2.7	112
4.2.10.15.1. Bug fixes	113

4.2.10.16. Mirror registry for Red Hat OpenShift 1.2.6	113
4.2.10.16.1. New features	113
4.2.10.17. Mirror registry for Red Hat OpenShift 1.2.5	113
4.2.10.18. Mirror registry for Red Hat OpenShift 1.2.4	113
4.2.10.19. Mirror registry for Red Hat OpenShift 1.2.3	113
4.2.10.20. Mirror registry for Red Hat OpenShift 1.2.2	113
4.2.10.21. Mirror registry for Red Hat OpenShift 1.2.1	113
4.2.10.22. Mirror registry for Red Hat OpenShift 1.2.0	114
4.2.10.22.1. Bug fixes	114
4.2.10.23. Mirror registry for Red Hat OpenShift 1.1.0	114
4.2.10.23.1. New features	114
4.2.10.23.2. Bug fixes	114
4.2.11. Troubleshooting mirror registry for Red Hat OpenShift	114
4.2.12. Additional resources	115
4.3. MIRRORING IMAGES FOR A DISCONNECTED INSTALLATION	115
4.3.1. Prerequisites	116
4.3.2. About the mirror registry	116
4.3.3. Preparing your mirror host	117
4.3.3.1. Installing the OpenShift CLI	117
Installing the OpenShift CLI on Linux	117
Installing the OpenShift CLI on Windows	118
Installing the OpenShift CLI on macOS	118
4.3.4. Configuring credentials that allow images to be mirrored	119
4.3.5. Mirroring the OpenShift Container Platform image repository	121
4.3.6. The Cluster Samples Operator in a disconnected environment	124
4.3.6.1. Cluster Samples Operator assistance for mirroring	124
4.3.7. Mirroring Operator catalogs for use with disconnected clusters	125
4.3.7.1. Prerequisites	126
4.3.7.2. Extracting and mirroring catalog contents	126
4.3.7.2.1. Mirroring catalog contents to registries on the same network	126
4.3.7.2.2. Mirroring catalog contents to airgapped registries	128
4.3.7.3. Generated manifests	130
4.3.7.4. Postinstallation requirements	131
4.3.8. Next steps	131
4.3.9. Additional resources	132
4.4. MIRRORING IMAGES FOR A DISCONNECTED INSTALLATION USING THE OC-MIRROR PLUGIN	132
4.4.1. About the oc-mirror plugin	132
4.4.1.1. High level workflow	132
4.4.2. oc-mirror plugin compatibility and support	133
4.4.3. About the mirror registry	133
4.4.4. Prerequisites	134
4.4.5. Preparing your mirror hosts	134
4.4.5.1. Installing the oc-mirror OpenShift CLI plugin	134
4.4.5.2. Configuring credentials that allow images to be mirrored	135
4.4.6. Creating the image set configuration	138
4.4.7. Mirroring an image set to a mirror registry	140
4.4.7.1. Mirroring an image set in a partially disconnected environment	140
4.4.7.1.1. Mirroring from mirror to mirror	140
4.4.7.2. Mirroring an image set in a fully disconnected environment	141
4.4.7.2.1. Mirroring from mirror to disk	141
4.4.7.2.2. Mirroring from disk to mirror	142
4.4.8. Configuring your cluster to use the resources generated by oc-mirror	143
4.4.9. Updating your mirror registry content	144

4.4.9.1. Mirror registry update examples	145
Mirroring a specific OpenShift Container Platform version by pruning the existing images	145
Updating to the latest version of an Operator by pruning the existing images	146
Mirroring a new Operator by pruning the existing Operator	146
Pruning all the OpenShift Container Platform images	147
4.4.10. Performing a dry run	147
4.4.11. Including local OCI Operator catalogs	148
4.4.12. Image set configuration parameters	151
4.4.13. Image set configuration examples	158
Use case: Including the shortest OpenShift Container Platform update path	158
Use case: Including all versions of OpenShift Container Platform from a minimum to the latest version for multi-architecture releases	158
Use case: Including Operator versions from a minimum to the latest	159
Use case: Including the Nutanix CSI Operator	160
Use case: Including the default Operator channel	160
Use case: Including an entire catalog (all versions)	161
Use case: Including an entire catalog (channel heads only)	161
Use case: Including arbitrary images and helm charts	161
4.4.14. Command reference for oc-mirror	162
4.4.15. Additional resources	164
4.5. MIRRORING IMAGES FOR A DISCONNECTED INSTALLATION BY USING THE OC-MIRROR PLUGIN V2	164
4.5.1. Prerequisites	164
4.5.2. About oc-mirror plugin v2	165
4.5.2.1. oc-mirror plugin v2 compatibility and support	166
4.5.3. Preparing your mirror hosts	166
4.5.3.1. Installing the oc-mirror OpenShift CLI plugin	166
4.5.3.2. Configuring credentials that allow images to be mirrored	167
4.5.4. Mirroring an image set to a mirror registry	169
4.5.4.1. Building the image set configuration	169
4.5.4.2. Mirroring an image set in a partially disconnected environment	170
4.5.4.3. Mirroring an image set in a fully disconnected environment	170
4.5.4.3.1. Mirroring from mirror to disk	170
4.5.4.3.2. Mirroring from disk to mirror	171
4.5.5. Additional resources	172
4.5.6. About custom resources generated by v2	172
4.5.6.1. Configuring your cluster to use the resources generated by oc-mirror plugin v2	172
4.5.7. Deletion of images from your disconnected environment	173
4.5.7.1. Deleting the images from disconnected environment	174
4.5.8. Verifying your selected images for mirroring	175
4.5.8.1. Performing dry run for oc-mirror plugin v2	175
4.5.9. Benefits of enclave support	176
4.5.9.1. Mirroring to an enclave	176
4.5.10. How filtering works in the operator catalog	179
4.5.11. ImageSet configuration parameters for oc-mirror plugin v2	184
4.5.11.1. Delete ImageSet Configuration parameters	190
4.5.12. Command reference for oc-mirror plugin v2	194
CHAPTER 5. INSTALLING ON ALIBABA CLOUD	197
5.1. INSTALLING A CLUSTER ON ALIBABA CLOUD BY USING THE ASSISTED INSTALLER	197
5.1.1. Process outline for creating a cluster with the Assisted Installer	197
5.1.2. Converting the discovery image to QCOW2 format	197

CHAPTER 6. INSTALLING ON AWS	199
6.1. INSTALLATION METHODS	199
6.1.1. Installing a cluster on installer-provisioned infrastructure	199
6.1.2. Installing a cluster on user-provisioned infrastructure	199
6.1.3. Installing a cluster on a single node	200
6.1.4. Additional resources	200
6.2. CONFIGURING AN AWS ACCOUNT	200
6.2.1. Configuring Route 53	200
6.2.1.1. Ingress Operator endpoint configuration for AWS Route 53	201
6.2.2. AWS account limits	202
6.2.3. Required AWS permissions for the IAM user	204
6.2.4. Creating an IAM user	213
6.2.5. IAM Policies and AWS authentication	213
6.2.5.1. Default permissions for IAM instance profiles	214
6.2.5.2. Specifying an existing IAM role	216
6.2.5.3. Using AWS IAM Analyzer to create policy templates	216
6.2.6. Supported AWS Marketplace regions	217
6.2.7. Supported AWS regions	217
6.2.7.1. AWS public regions	218
6.2.7.2. AWS GovCloud regions	219
6.2.7.3. AWS SC2S and C2S secret regions	219
6.2.7.4. AWS China regions	219
6.2.8. Next steps	219
6.3. INSTALLER-PROVISIONED INFRASTRUCTURE	219
6.3.1. Preparing to install a cluster on AWS	219
6.3.1.1. Internet access for OpenShift Container Platform	220
6.3.1.2. Obtaining the installation program	220
6.3.1.3. Installing the OpenShift CLI	221
Installing the OpenShift CLI on Linux	221
Installing the OpenShift CLI on Windows	222
Installing the OpenShift CLI on macOS	222
6.3.1.4. Generating a key pair for cluster node SSH access	223
6.3.1.5. Telemetry access for OpenShift Container Platform	225
6.3.2. Installing a cluster on AWS	225
6.3.2.1. Prerequisites	225
6.3.2.2. Deploying the cluster	225
6.3.2.3. Logging in to the cluster by using the CLI	228
6.3.2.4. Logging in to the cluster by using the web console	228
6.3.2.5. Next steps	229
6.3.3. Installing a cluster on AWS with customizations	230
6.3.3.1. Prerequisites	230
6.3.3.2. Obtaining an AWS Marketplace image	230
6.3.3.3. Creating the installation configuration file	231
6.3.3.3.1. Minimum resource requirements for cluster installation	233
6.3.3.3.2. Tested instance types for AWS	234
6.3.3.3.3. Tested instance types for AWS on 64-bit ARM infrastructures	235
6.3.3.3.4. Sample customized install-config.yaml file for AWS	236
6.3.3.3.5. Configuring the cluster-wide proxy during installation	239
6.3.3.4. Alternatives to storing administrator-level secrets in the kube-system project	241
6.3.3.4.1. Manually creating long-term credentials	241
6.3.3.4.2. Configuring an AWS cluster to use short-term credentials	243
6.3.3.4.2.1. Configuring the Cloud Credential Operator utility	243
6.3.3.4.2.2. Creating AWS resources with the Cloud Credential Operator utility	246

6.3.3.4.2.2.1. Creating AWS resources with a single command	247
6.3.3.4.2.2.2. Creating AWS resources individually	249
6.3.3.4.2.3. Incorporating the Cloud Credential Operator utility manifests	251
6.3.3.5. Deploying the cluster	252
6.3.3.6. Logging in to the cluster by using the CLI	254
6.3.3.7. Logging in to the cluster by using the web console	254
6.3.3.8. Next steps	255
6.3.4. Installing a cluster on AWS with network customizations	255
6.3.4.1. Prerequisites	256
6.3.4.2. Network configuration phases	256
6.3.4.3. Creating the installation configuration file	257
6.3.4.3.1. Minimum resource requirements for cluster installation	259
6.3.4.3.2. Tested instance types for AWS	260
6.3.4.3.3. Tested instance types for AWS on 64-bit ARM infrastructures	261
6.3.4.3.4. Sample customized install-config.yaml file for AWS	261
6.3.4.3.5. Configuring the cluster-wide proxy during installation	265
6.3.4.4. Alternatives to storing administrator-level secrets in the kube-system project	266
6.3.4.4.1. Manually creating long-term credentials	266
6.3.4.4.2. Configuring an AWS cluster to use short-term credentials	269
6.3.4.4.2.1. Configuring the Cloud Credential Operator utility	269
6.3.4.4.2.2. Creating AWS resources with the Cloud Credential Operator utility	272
6.3.4.4.2.2.1. Creating AWS resources with a single command	272
6.3.4.4.2.2.2. Creating AWS resources individually	274
6.3.4.4.2.3. Incorporating the Cloud Credential Operator utility manifests	277
6.3.4.5. Cluster Network Operator configuration	278
6.3.4.5.1. Cluster Network Operator configuration object	278
defaultNetwork object configuration	279
Configuration for the OVN-Kubernetes network plugin	280
kubeProxyConfig object configuration (OpenShiftSDN container network interface only)	284
6.3.4.6. Specifying advanced network configuration	285
6.3.4.7. Configuring an Ingress Controller Network Load Balancer on a new AWS cluster	286
6.3.4.8. Configuring hybrid networking with OVN-Kubernetes	287
6.3.4.9. Deploying the cluster	288
6.3.4.10. Logging in to the cluster by using the CLI	290
6.3.4.11. Logging in to the cluster by using the web console	290
6.3.4.12. Next steps	291
6.3.5. Installing a cluster on AWS in a restricted network	291
6.3.5.1. Prerequisites	292
6.3.5.2. About installations in restricted networks	292
6.3.5.2.1. Additional limits	293
6.3.5.3. About using a custom VPC	293
6.3.5.3.1. Requirements for using your VPC	293
Option 1: Create VPC endpoints	294
Option 2: Create a proxy without VPC endpoints	295
Option 3: Create a proxy with VPC endpoints	295
6.3.5.3.2. VPC validation	296
6.3.5.3.3. Division of permissions	297
6.3.5.3.4. Isolation between clusters	297
6.3.5.4. Creating the installation configuration file	297
6.3.5.4.1. Minimum resource requirements for cluster installation	300
6.3.5.4.2. Sample customized install-config.yaml file for AWS	301
6.3.5.4.3. Configuring the cluster-wide proxy during installation	304
6.3.5.5. Alternatives to storing administrator-level secrets in the kube-system project	306

6.3.5.5.1. Manually creating long-term credentials	306
6.3.5.5.2. Configuring an AWS cluster to use short-term credentials	308
6.3.5.5.2.1. Configuring the Cloud Credential Operator utility	308
6.3.5.5.2.2. Creating AWS resources with the Cloud Credential Operator utility	312
6.3.5.5.2.2.1. Creating AWS resources with a single command	312
6.3.5.5.2.2.2. Creating AWS resources individually	314
6.3.5.5.2.3. Incorporating the Cloud Credential Operator utility manifests	317
6.3.5.6. Deploying the cluster	318
6.3.5.7. Logging in to the cluster by using the CLI	319
6.3.5.8. Disabling the default OperatorHub catalog sources	320
6.3.5.9. Next steps	320
6.3.6. Installing a cluster on AWS into an existing VPC	320
6.3.6.1. Prerequisites	320
6.3.6.2. About using a custom VPC	321
6.3.6.2.1. Requirements for using your VPC	321
Option 1: Create VPC endpoints	323
Option 2: Create a proxy without VPC endpoints	323
Option 3: Create a proxy with VPC endpoints	323
6.3.6.2.2. VPC validation	325
6.3.6.2.3. Division of permissions	325
6.3.6.2.4. Isolation between clusters	325
6.3.6.2.5. Optional: AWS security groups	325
6.3.6.2.6. Modifying trust policy when installing into a shared VPC	326
6.3.6.3. Creating the installation configuration file	326
6.3.6.3.1. Minimum resource requirements for cluster installation	328
6.3.6.3.2. Tested instance types for AWS	329
6.3.6.3.3. Tested instance types for AWS on 64-bit ARM infrastructures	330
6.3.6.3.4. Sample customized install-config.yaml file for AWS	330
6.3.6.3.5. Configuring the cluster-wide proxy during installation	333
6.3.6.3.6. Applying existing AWS security groups to the cluster	335
6.3.6.4. Alternatives to storing administrator-level secrets in the kube-system project	336
6.3.6.4.1. Manually creating long-term credentials	336
6.3.6.4.2. Configuring an AWS cluster to use short-term credentials	338
6.3.6.4.2.1. Configuring the Cloud Credential Operator utility	338
6.3.6.4.2.2. Creating AWS resources with the Cloud Credential Operator utility	342
6.3.6.4.2.2.1. Creating AWS resources with a single command	342
6.3.6.4.2.2.2. Creating AWS resources individually	344
6.3.6.4.2.3. Incorporating the Cloud Credential Operator utility manifests	347
6.3.6.5. Deploying the cluster	347
6.3.6.6. Logging in to the cluster by using the CLI	349
6.3.6.7. Logging in to the cluster by using the web console	350
6.3.6.8. Next steps	350
6.3.7. Installing a private cluster on AWS	351
6.3.7.1. Prerequisites	351
6.3.7.2. Private clusters	351
6.3.7.2.1. Private clusters in AWS	352
6.3.7.2.1.1. Limitations	352
6.3.7.3. About using a custom VPC	353
6.3.7.3.1. Requirements for using your VPC	353
Option 1: Create VPC endpoints	354
Option 2: Create a proxy without VPC endpoints	354
Option 3: Create a proxy with VPC endpoints	354
6.3.7.3.2. VPC validation	356

6.3.7.3.3. Division of permissions	356
6.3.7.3.4. Isolation between clusters	356
6.3.7.3.5. Optional: AWS security groups	357
6.3.7.4. Manually creating the installation configuration file	357
6.3.7.4.1. Minimum resource requirements for cluster installation	358
6.3.7.4.2. Tested instance types for AWS	359
6.3.7.4.3. Tested instance types for AWS on 64-bit ARM infrastructures	360
6.3.7.4.4. Sample customized install-config.yaml file for AWS	360
6.3.7.4.5. Configuring the cluster-wide proxy during installation	363
6.3.7.4.6. Applying existing AWS security groups to the cluster	365
6.3.7.5. Alternatives to storing administrator-level secrets in the kube-system project	366
6.3.7.5.1. Manually creating long-term credentials	366
6.3.7.5.2. Configuring an AWS cluster to use short-term credentials	368
6.3.7.5.2.1. Configuring the Cloud Credential Operator utility	368
6.3.7.5.2.2. Creating AWS resources with the Cloud Credential Operator utility	372
6.3.7.5.2.2.1. Creating AWS resources with a single command	372
6.3.7.5.2.2.2. Creating AWS resources individually	374
6.3.7.5.2.3. Incorporating the Cloud Credential Operator utility manifests	377
6.3.7.6. Deploying the cluster	377
6.3.7.7. Logging in to the cluster by using the CLI	379
6.3.7.8. Logging in to the cluster by using the web console	380
6.3.7.9. Next steps	380
6.3.8. Installing a cluster on AWS into a government region	381
6.3.8.1. Prerequisites	381
6.3.8.2. AWS government regions	381
6.3.8.3. Installation requirements	381
6.3.8.4. Private clusters	382
6.3.8.4.1. Private clusters in AWS	382
6.3.8.4.1.1. Limitations	383
6.3.8.5. About using a custom VPC	383
6.3.8.5.1. Requirements for using your VPC	383
Option 1: Create VPC endpoints	384
Option 2: Create a proxy without VPC endpoints	385
Option 3: Create a proxy with VPC endpoints	385
6.3.8.5.2. VPC validation	386
6.3.8.5.3. Division of permissions	387
6.3.8.5.4. Isolation between clusters	387
6.3.8.5.5. Optional: AWS security groups	387
6.3.8.6. Obtaining an AWS Marketplace image	387
6.3.8.7. Manually creating the installation configuration file	388
6.3.8.7.1. Minimum resource requirements for cluster installation	389
6.3.8.7.2. Tested instance types for AWS	390
6.3.8.7.3. Tested instance types for AWS on 64-bit ARM infrastructures	391
6.3.8.7.4. Sample customized install-config.yaml file for AWS	391
6.3.8.7.5. Configuring the cluster-wide proxy during installation	394
6.3.8.7.6. Applying existing AWS security groups to the cluster	396
6.3.8.8. Alternatives to storing administrator-level secrets in the kube-system project	397
6.3.8.8.1. Manually creating long-term credentials	397
6.3.8.8.2. Configuring an AWS cluster to use short-term credentials	399
6.3.8.8.2.1. Configuring the Cloud Credential Operator utility	400
6.3.8.8.2.2. Creating AWS resources with the Cloud Credential Operator utility	403
6.3.8.8.2.2.1. Creating AWS resources with a single command	403
6.3.8.8.2.2.2. Creating AWS resources individually	405

6.3.8.8.2.3. Incorporating the Cloud Credential Operator utility manifests	408
6.3.8.9. Deploying the cluster	409
6.3.8.10. Logging in to the cluster by using the CLI	410
6.3.8.11. Logging in to the cluster by using the web console	411
6.3.8.12. Next steps	412
6.3.9. Installing a cluster on AWS into a Secret or Top Secret Region	412
6.3.9.1. Prerequisites	412
6.3.9.2. AWS secret regions	413
6.3.9.3. Installation requirements	413
6.3.9.4. Private clusters	414
6.3.9.4.1. Private clusters in AWS	414
6.3.9.4.1.1. Limitations	415
6.3.9.5. About using a custom VPC	415
6.3.9.5.1. Requirements for using your VPC	415
Option 1: Create VPC endpoints	416
Option 2: Create a proxy without VPC endpoints	417
Option 3: Create a proxy with VPC endpoints	417
6.3.9.5.2. VPC validation	419
6.3.9.5.3. Division of permissions	419
6.3.9.5.4. Isolation between clusters	419
6.3.9.5.5. Optional: AWS security groups	419
6.3.9.6. Uploading a custom RHCOS AMI in AWS	420
6.3.9.7. Manually creating the installation configuration file	422
6.3.9.7.1. Tested instance types for AWS	423
6.3.9.7.2. Sample customized install-config.yaml file for AWS	423
6.3.9.7.3. Configuring the cluster-wide proxy during installation	427
6.3.9.7.4. Applying existing AWS security groups to the cluster	428
6.3.9.8. Alternatives to storing administrator-level secrets in the kube-system project	429
6.3.9.8.1. Manually creating long-term credentials	430
6.3.9.8.2. Configuring an AWS cluster to use short-term credentials	432
6.3.9.8.2.1. Configuring the Cloud Credential Operator utility	432
6.3.9.8.2.2. Creating AWS resources with the Cloud Credential Operator utility	435
6.3.9.8.2.2.1. Creating AWS resources with a single command	435
6.3.9.8.2.2.2. Creating AWS resources individually	437
6.3.9.8.2.3. Incorporating the Cloud Credential Operator utility manifests	440
6.3.9.9. Deploying the cluster	441
6.3.9.10. Logging in to the cluster by using the CLI	443
6.3.9.11. Logging in to the cluster by using the web console	443
6.3.9.12. Next steps	444
6.3.10. Installing a cluster on AWS China	444
6.3.10.1. Prerequisites	444
6.3.10.2. Installation requirements	445
6.3.10.3. Private clusters	445
6.3.10.3.1. Private clusters in AWS	446
6.3.10.3.1.1. Limitations	446
6.3.10.4. About using a custom VPC	446
6.3.10.4.1. Requirements for using your VPC	447
Option 1: Create VPC endpoints	448
Option 2: Create a proxy without VPC endpoints	448
Option 3: Create a proxy with VPC endpoints	448
6.3.10.4.2. VPC validation	450
6.3.10.4.3. Division of permissions	450
6.3.10.4.4. Isolation between clusters	450

6.3.10.4.5. Optional: AWS security groups	450
6.3.10.5. Uploading a custom RHCOS AMI in AWS	451
6.3.10.6. Manually creating the installation configuration file	453
6.3.10.6.1. Sample customized install-config.yaml file for AWS	454
6.3.10.6.2. Minimum resource requirements for cluster installation	457
6.3.10.6.3. Tested instance types for AWS	458
6.3.10.6.4. Tested instance types for AWS on 64-bit ARM infrastructures	459
6.3.10.6.5. Configuring the cluster-wide proxy during installation	459
6.3.10.6.6. Applying existing AWS security groups to the cluster	461
6.3.10.7. Alternatives to storing administrator-level secrets in the kube-system project	462
6.3.10.7.1. Manually creating long-term credentials	462
6.3.10.7.2. Configuring an AWS cluster to use short-term credentials	464
6.3.10.7.2.1. Configuring the Cloud Credential Operator utility	465
6.3.10.7.2.2. Creating AWS resources with the Cloud Credential Operator utility	468
6.3.10.7.2.2.1. Creating AWS resources with a single command	468
6.3.10.7.2.2.2. Creating AWS resources individually	470
6.3.10.7.2.3. Incorporating the Cloud Credential Operator utility manifests	473
6.3.10.8. Deploying the cluster	474
6.3.10.9. Logging in to the cluster by using the CLI	475
6.3.10.10. Logging in to the cluster by using the web console	476
6.3.10.11. Next steps	477
6.3.11. Installing a cluster with compute nodes on AWS Local Zones	477
6.3.11.1. Infrastructure prerequisites	477
6.3.11.2. About AWS Local Zones and edge compute pool	478
6.3.11.2.1. Cluster limitations in AWS Local Zones	478
6.3.11.2.2. About edge compute pools	479
6.3.11.3. Installation prerequisites	480
6.3.11.3.1. Opting in to an AWS Local Zones	480
6.3.11.3.2. Obtaining an AWS Marketplace image	481
6.3.11.4. Preparing for the installation	482
6.3.11.4.1. Minimum resource requirements for cluster installation	482
6.3.11.4.2. Tested instance types for AWS	483
6.3.11.4.3. Creating the installation configuration file	484
6.3.11.4.4. Examples of installation configuration files with edge compute pools	485
6.3.11.4.5. Customizing the cluster network MTU	487
6.3.11.5. Cluster installation options for an AWS Local Zones environment	488
6.3.11.6. Install a cluster quickly in AWS Local Zones	488
6.3.11.6.1. Modifying an installation configuration file to use AWS Local Zones	488
6.3.11.7. Installing a cluster in an existing VPC that has Local Zone subnets	490
6.3.11.7.1. Creating a VPC in AWS	490
6.3.11.7.2. CloudFormation template for the VPC	492
6.3.11.7.3. Creating subnets in Local Zones	498
6.3.11.7.4. CloudFormation template for the VPC subnet	500
6.3.11.7.5. Modifying an installation configuration file to use AWS Local Zones subnets	502
6.3.11.8. Optional: AWS security groups	502
6.3.11.9. Optional: Assign public IP addresses to edge compute nodes	503
6.3.11.10. Deploying the cluster	504
6.3.11.11. Verifying the status of the deployed cluster	505
6.3.11.11.1. Logging in to the cluster by using the CLI	506
6.3.11.11.2. Logging in to the cluster by using the web console	506
6.3.11.11.3. Verifying nodes that were created with edge compute pool	507
6.3.12. Installing a cluster with compute nodes on AWS Wavelength Zones	508
6.3.12.1. Infrastructure prerequisites	508

6.3.12.2. About AWS Wavelength Zones and edge compute pool	509
6.3.12.2.1. Cluster limitations in AWS Wavelength Zones	510
6.3.12.2.2. About edge compute pools	510
6.3.12.3. Installation prerequisites	511
6.3.12.3.1. Opting in to an AWS Wavelength Zones	512
6.3.12.3.2. Obtaining an AWS Marketplace image	512
6.3.12.4. Preparing for the installation	513
6.3.12.4.1. Minimum resource requirements for cluster installation	513
6.3.12.4.2. Tested instance types for AWS	514
6.3.12.4.3. Creating the installation configuration file	515
6.3.12.4.4. Examples of installation configuration files with edge compute pools	516
6.3.12.5. Cluster installation options for an AWS Wavelength Zones environment	517
6.3.12.6. Install a cluster quickly in AWS Wavelength Zones	518
6.3.12.6.1. Modifying an installation configuration file to use AWS Wavelength Zones	518
6.3.12.7. Installing a cluster in an existing VPC that has Wavelength Zone subnets	519
6.3.12.7.1. Creating a VPC in AWS	520
6.3.12.7.2. CloudFormation template for the VPC	522
6.3.12.7.3. Creating a VPC carrier gateway	528
6.3.12.7.4. CloudFormation template for the VPC Carrier Gateway	530
6.3.12.7.5. Creating subnets in Wavelength Zones	531
6.3.12.7.6. CloudFormation template for the VPC subnet	533
6.3.12.7.7. Modifying an installation configuration file to use AWS Wavelength Zones subnets	535
6.3.12.8. Optional: Assign public IP addresses to edge compute nodes	535
6.3.12.9. Deploying the cluster	537
6.3.12.10. Verifying the status of the deployed cluster	538
6.3.12.10.1. Logging in to the cluster by using the CLI	538
6.3.12.10.2. Logging in to the cluster by using the web console	539
6.3.12.10.3. Verifying nodes that were created with edge compute pool	540
6.3.13. Installing a cluster with compute nodes on AWS Outposts	541
6.4. USER-PROVISIONED INFRASTRUCTURE	541
6.4.1. Preparing to install a cluster on AWS	541
6.4.1.1. Internet access for OpenShift Container Platform	542
6.4.1.2. Obtaining the installation program	542
6.4.1.3. Installing the OpenShift CLI	543
Installing the OpenShift CLI on Linux	543
Installing the OpenShift CLI on Windows	544
Installing the OpenShift CLI on macOS	544
6.4.1.4. Generating a key pair for cluster node SSH access	545
6.4.1.5. Telemetry access for OpenShift Container Platform	546
6.4.2. Installation requirements for user-provisioned infrastructure on AWS	547
6.4.2.1. Required machines for cluster installation	547
6.4.2.1.1. Minimum resource requirements for cluster installation	547
6.4.2.1.2. Tested instance types for AWS	548
6.4.2.1.3. Tested instance types for AWS on 64-bit ARM infrastructures	549
6.4.2.2. Certificate signing requests management	550
6.4.2.3. Required AWS infrastructure components	550
6.4.2.3.1. Other infrastructure components	550
Option 1: Create VPC endpoints	551
Option 2: Create a proxy without VPC endpoints	551
Option 3: Create a proxy with VPC endpoints	551
6.4.2.3.2. Cluster machines	558
6.4.2.4. Required AWS permissions for the IAM user	559
6.4.2.5. Obtaining an AWS Marketplace image	567

6.4.3. Installing a cluster on user-provisioned infrastructure in AWS by using CloudFormation templates	567
6.4.3.1. Prerequisites	568
6.4.3.2. Creating the installation files for AWS	568
6.4.3.2.1. Optional: Creating a separate /var partition	569
6.4.3.2.2. Creating the installation configuration file	571
6.4.3.2.3. Configuring the cluster-wide proxy during installation	573
6.4.3.2.4. Creating the Kubernetes manifest and Ignition config files	574
6.4.3.3. Extracting the infrastructure name	577
6.4.3.4. Creating a VPC in AWS	577
6.4.3.4.1. CloudFormation template for the VPC	579
6.4.3.5. Creating networking and load balancing components in AWS	585
6.4.3.5.1. CloudFormation template for the network and load balancers	589
6.4.3.6. Creating security group and roles in AWS	597
6.4.3.6.1. CloudFormation template for security objects	599
6.4.3.7. Accessing RHCOS AMIs with stream metadata	610
6.4.3.8. RHCOS AMIs for the AWS infrastructure	611
6.4.3.8.1. AWS regions without a published RHCOS AMI	614
6.4.3.8.2. Uploading a custom RHCOS AMI in AWS	615
6.4.3.9. Creating the bootstrap node in AWS	617
6.4.3.9.1. CloudFormation template for the bootstrap machine	622
6.4.3.10. Creating the control plane machines in AWS	626
6.4.3.10.1. CloudFormation template for control plane machines	630
6.4.3.11. Creating the worker nodes in AWS	636
6.4.3.11.1. CloudFormation template for worker machines	639
6.4.3.12. Initializing the bootstrap sequence on AWS with user-provisioned infrastructure	642
6.4.3.13. Logging in to the cluster by using the CLI	643
6.4.3.14. Approving the certificate signing requests for your machines	644
6.4.3.15. Initial Operator configuration	646
6.4.3.15.1. Image registry storage configuration	647
6.4.3.15.1.1. Configuring registry storage for AWS with user-provisioned infrastructure	648
6.4.3.15.1.2. Configuring storage for the image registry in non-production clusters	648
6.4.3.16. Deleting the bootstrap resources	649
6.4.3.17. Creating the Ingress DNS Records	649
6.4.3.18. Completing an AWS installation on user-provisioned infrastructure	652
6.4.3.19. Logging in to the cluster by using the web console	653
6.4.3.20. Additional resources	654
6.4.3.21. Next steps	654
6.4.4. Installing a cluster on AWS in a restricted network with user-provisioned infrastructure	654
6.4.4.1. Prerequisites	655
6.4.4.2. About installations in restricted networks	656
6.4.4.2.1. Additional limits	656
6.4.4.3. Creating the installation files for AWS	656
6.4.4.3.1. Optional: Creating a separate /var partition	656
6.4.4.3.2. Creating the installation configuration file	659
6.4.4.3.3. Configuring the cluster-wide proxy during installation	661
6.4.4.3.4. Creating the Kubernetes manifest and Ignition config files	663
6.4.4.4. Extracting the infrastructure name	665
6.4.4.5. Creating a VPC in AWS	666
6.4.4.5.1. CloudFormation template for the VPC	667
6.4.4.6. Creating networking and load balancing components in AWS	673
6.4.4.6.1. CloudFormation template for the network and load balancers	677
6.4.4.7. Creating security group and roles in AWS	685
6.4.4.7.1. CloudFormation template for security objects	687

6.4.4.8. Accessing RHCOS AMIs with stream metadata	698
6.4.4.9. RHCOS AMIs for the AWS infrastructure	699
6.4.4.10. Creating the bootstrap node in AWS	702
6.4.4.10.1. CloudFormation template for the bootstrap machine	707
6.4.4.11. Creating the control plane machines in AWS	711
6.4.4.11.1. CloudFormation template for control plane machines	716
6.4.4.12. Creating the worker nodes in AWS	721
6.4.4.12.1. CloudFormation template for worker machines	725
6.4.4.13. Initializing the bootstrap sequence on AWS with user-provisioned infrastructure	727
6.4.4.14. Approving the certificate signing requests for your machines	728
6.4.4.15. Initial Operator configuration	731
6.4.4.15.1. Disabling the default OperatorHub catalog sources	732
6.4.4.15.2. Image registry storage configuration	732
6.4.4.15.2.1. Configuring registry storage for AWS with user-provisioned infrastructure	732
6.4.4.15.2.2. Configuring storage for the image registry in non-production clusters	733
6.4.4.16. Deleting the bootstrap resources	734
6.4.4.17. Creating the Ingress DNS Records	734
6.4.4.18. Completing an AWS installation on user-provisioned infrastructure	737
6.4.4.19. Logging in to the cluster by using the CLI	738
6.4.4.20. Logging in to the cluster by using the web console	738
6.4.4.21. Additional resources	739
6.4.4.22. Next steps	739
6.5. INSTALLING A THREE-NODE CLUSTER ON AWS	740
6.5.1. Configuring a three-node cluster	740
6.5.2. Next steps	741
6.6. UNINSTALLING A CLUSTER ON AWS	741
6.6.1. Removing a cluster that uses installer-provisioned infrastructure	741
6.6.2. Deleting Amazon Web Services resources with the Cloud Credential Operator utility	742
6.6.3. Deleting a cluster with a configured AWS Local Zone infrastructure	743
6.7. INSTALLATION CONFIGURATION PARAMETERS FOR AWS	744
6.7.1. Available installation configuration parameters for AWS	744
6.7.1.1. Required configuration parameters	745
6.7.1.2. Network configuration parameters	746
6.7.1.3. Optional configuration parameters	747
6.7.1.4. Optional AWS configuration parameters	754
CHAPTER 7. INSTALLING ON AZURE	761
7.1. PREPARING TO INSTALL ON AZURE	761
7.1.1. Prerequisites	761
7.1.2. Requirements for installing OpenShift Container Platform on Azure	761
7.1.3. Choosing a method to install OpenShift Container Platform on Azure	761
7.1.3.1. Installing a cluster on installer-provisioned infrastructure	761
7.1.3.2. Installing a cluster on user-provisioned infrastructure	762
7.1.4. Next steps	762
7.2. CONFIGURING AN AZURE ACCOUNT	762
7.2.1. Azure account limits	762
7.2.2. Configuring a public DNS zone in Azure	765
7.2.3. Increasing Azure account limits	765
7.2.4. Recording the subscription and tenant IDs	766
7.2.5. Supported identities to access Azure resources	768
7.2.5.1. Required Azure roles	768
7.2.5.2. Required Azure permissions for installer-provisioned infrastructure	768
7.2.5.3. Using Azure managed identities	776

7.2.5.4. Creating a service principal	776
7.2.6. Supported Azure Marketplace regions	777
7.2.7. Supported Azure regions	778
Supported Azure public regions	778
Supported Azure Government regions	779
7.2.8. Next steps	779
7.3. ENABLING USER-MANAGED ENCRYPTION FOR AZURE	779
7.3.1. Preparing an Azure Disk Encryption Set	780
7.3.2. Next steps	782
7.4. INSTALLING A CLUSTER QUICKLY ON AZURE	782
7.4.1. Prerequisites	782
7.4.2. Internet access for OpenShift Container Platform	782
7.4.3. Generating a key pair for cluster node SSH access	783
7.4.4. Obtaining the installation program	784
7.4.5. Deploying the cluster	785
7.4.6. Installing the OpenShift CLI	788
Installing the OpenShift CLI on Linux	788
Installing the OpenShift CLI on Windows	789
Installing the OpenShift CLI on macOS	789
7.4.7. Logging in to the cluster by using the CLI	790
7.4.8. Telemetry access for OpenShift Container Platform	790
7.4.9. Next steps	791
7.5. INSTALLING A CLUSTER ON AZURE WITH CUSTOMIZATIONS	791
7.5.1. Prerequisites	791
7.5.2. Internet access for OpenShift Container Platform	791
7.5.3. Generating a key pair for cluster node SSH access	791
7.5.4. Using the Azure Marketplace offering	793
7.5.5. Obtaining the installation program	796
7.5.6. Creating the installation configuration file	796
7.5.6.1. Minimum resource requirements for cluster installation	799
7.5.6.2. Tested instance types for Azure	800
7.5.6.3. Tested instance types for Azure on 64-bit ARM infrastructures	801
7.5.6.4. Enabling trusted launch for Azure VMs	801
7.5.6.5. Enabling confidential VMs	802
7.5.6.6. Sample customized install-config.yaml file for Azure	803
7.5.6.7. Configuring the cluster-wide proxy during installation	806
7.5.7. Configuring user-defined tags for Azure	808
7.5.8. Querying user-defined tags for Azure	810
7.5.9. Installing the OpenShift CLI	810
Installing the OpenShift CLI on Linux	811
Installing the OpenShift CLI on Windows	811
Installing the OpenShift CLI on macOS	812
7.5.10. Alternatives to storing administrator-level secrets in the kube-system project	812
7.5.10.1. Manually creating long-term credentials	812
7.5.10.2. Configuring an Azure cluster to use short-term credentials	814
7.5.10.2.1. Configuring the Cloud Credential Operator utility	815
7.5.10.2.2. Creating Azure resources with the Cloud Credential Operator utility	817
7.5.10.2.3. Incorporating the Cloud Credential Operator utility manifests	819
7.5.11. Deploying the cluster	820
7.5.12. Logging in to the cluster by using the CLI	822
7.5.13. Telemetry access for OpenShift Container Platform	822
7.5.14. Next steps	823
7.6. INSTALLING A CLUSTER ON AZURE WITH NETWORK CUSTOMIZATIONS	823

7.6.1. Prerequisites	823
7.6.2. Internet access for OpenShift Container Platform	823
7.6.3. Generating a key pair for cluster node SSH access	824
7.6.4. Obtaining the installation program	826
7.6.5. Creating the installation configuration file	826
7.6.5.1. Minimum resource requirements for cluster installation	829
7.6.5.2. Tested instance types for Azure	830
7.6.5.3. Tested instance types for Azure on 64-bit ARM infrastructures	831
7.6.5.4. Enabling trusted launch for Azure VMs	831
7.6.5.5. Enabling confidential VMs	832
7.6.5.6. Sample customized install-config.yaml file for Azure	833
7.6.5.7. Configuring the cluster-wide proxy during installation	836
7.6.6. Network configuration phases	838
7.6.7. Specifying advanced network configuration	838
7.6.8. Cluster Network Operator configuration	839
7.6.8.1. Cluster Network Operator configuration object	840
defaultNetwork object configuration	840
Configuration for the OVN-Kubernetes network plugin	841
kubeProxyConfig object configuration (OpenShiftSDN container network interface only)	845
7.6.9. Configuring hybrid networking with OVN-Kubernetes	846
7.6.10. Installing the OpenShift CLI	847
Installing the OpenShift CLI on Linux	847
Installing the OpenShift CLI on Windows	848
Installing the OpenShift CLI on macOS	848
7.6.11. Alternatives to storing administrator-level secrets in the kube-system project	849
7.6.11.1. Manually creating long-term credentials	849
7.6.11.2. Configuring an Azure cluster to use short-term credentials	851
7.6.11.2.1. Configuring the Cloud Credential Operator utility	851
7.6.11.2.2. Creating Azure resources with the Cloud Credential Operator utility	854
7.6.11.2.3. Incorporating the Cloud Credential Operator utility manifests	856
7.6.12. Deploying the cluster	857
7.6.13. Logging in to the cluster by using the CLI	858
7.6.14. Telemetry access for OpenShift Container Platform	859
7.6.15. Next steps	859
7.7. INSTALLING A CLUSTER ON AZURE INTO AN EXISTING VNET	859
7.7.1. Prerequisites	859
7.7.2. About reusing a VNet for your OpenShift Container Platform cluster	860
7.7.2.1. Requirements for using your VNet	860
7.7.2.1.1. Network security group requirements	861
7.7.2.2. Division of permissions	862
7.7.2.3. Isolation between clusters	862
7.7.3. Internet access for OpenShift Container Platform	863
7.7.4. Generating a key pair for cluster node SSH access	863
7.7.5. Obtaining the installation program	865
7.7.6. Creating the installation configuration file	866
7.7.6.1. Minimum resource requirements for cluster installation	868
7.7.6.2. Tested instance types for Azure	869
7.7.6.3. Tested instance types for Azure on 64-bit ARM infrastructures	870
7.7.6.4. Enabling trusted launch for Azure VMs	870
7.7.6.5. Enabling confidential VMs	871
7.7.6.6. Sample customized install-config.yaml file for Azure	872
7.7.6.7. Configuring the cluster-wide proxy during installation	875
7.7.7. Installing the OpenShift CLI	877

Installing the OpenShift CLI on Linux	877
Installing the OpenShift CLI on Windows	878
Installing the OpenShift CLI on macOS	878
7.7.8. Alternatives to storing administrator-level secrets in the kube-system project	879
7.7.8.1. Manually creating long-term credentials	879
7.7.8.2. Configuring an Azure cluster to use short-term credentials	881
7.7.8.2.1. Configuring the Cloud Credential Operator utility	881
7.7.8.2.2. Creating Azure resources with the Cloud Credential Operator utility	883
7.7.8.2.3. Incorporating the Cloud Credential Operator utility manifests	886
7.7.9. Deploying the cluster	887
7.7.10. Telemetry access for OpenShift Container Platform	888
7.7.11. Next steps	888
7.8. INSTALLING A PRIVATE CLUSTER ON AZURE	889
7.8.1. Prerequisites	889
7.8.2. Private clusters	889
7.8.2.1. Private clusters in Azure	890
7.8.2.1.1. Limitations	890
7.8.2.2. User-defined outbound routing	890
Private cluster with network address translation	891
Private cluster with Azure Firewall	891
Private cluster with a proxy configuration	891
Private cluster with no internet access	891
7.8.3. About reusing a VNet for your OpenShift Container Platform cluster	891
7.8.3.1. Requirements for using your VNet	892
7.8.3.1.1. Network security group requirements	893
7.8.3.2. Division of permissions	894
7.8.3.3. Isolation between clusters	894
7.8.4. Internet access for OpenShift Container Platform	894
7.8.5. Generating a key pair for cluster node SSH access	894
7.8.6. Obtaining the installation program	896
7.8.7. Manually creating the installation configuration file	897
7.8.7.1. Minimum resource requirements for cluster installation	898
7.8.7.2. Tested instance types for Azure	899
7.8.7.3. Tested instance types for Azure on 64-bit ARM infrastructures	900
7.8.7.4. Enabling trusted launch for Azure VMs	900
7.8.7.5. Enabling confidential VMs	901
7.8.7.6. Sample customized install-config.yaml file for Azure	902
7.8.7.7. Configuring the cluster-wide proxy during installation	905
7.8.8. Installing the OpenShift CLI	907
Installing the OpenShift CLI on Linux	907
Installing the OpenShift CLI on Windows	908
Installing the OpenShift CLI on macOS	908
7.8.9. Alternatives to storing administrator-level secrets in the kube-system project	909
7.8.9.1. Manually creating long-term credentials	909
7.8.9.2. Configuring an Azure cluster to use short-term credentials	911
7.8.9.2.1. Configuring the Cloud Credential Operator utility	911
7.8.9.2.2. Creating Azure resources with the Cloud Credential Operator utility	914
7.8.9.2.3. Incorporating the Cloud Credential Operator utility manifests	916
7.8.10. Optional: Preparing a private Microsoft Azure cluster for a private image registry	917
7.8.11. Deploying the cluster	919
7.8.12. Logging in to the cluster by using the CLI	921
7.8.13. Telemetry access for OpenShift Container Platform	921
7.8.14. Next steps	922

7.9. INSTALLING A CLUSTER ON AZURE INTO A GOVERNMENT REGION	922
7.9.1. Prerequisites	922
7.9.2. Azure government regions	922
7.9.3. Private clusters	923
7.9.3.1. Private clusters in Azure	923
7.9.3.1.1. Limitations	924
7.9.3.2. User-defined outbound routing	924
7.9.4. About reusing a VNet for your OpenShift Container Platform cluster	924
7.9.4.1. Requirements for using your VNet	925
7.9.4.1.1. Network security group requirements	926
7.9.4.2. Division of permissions	927
7.9.4.3. Isolation between clusters	927
7.9.5. Internet access for OpenShift Container Platform	927
7.9.6. Generating a key pair for cluster node SSH access	928
7.9.7. Obtaining the installation program	930
7.9.8. Manually creating the installation configuration file	930
7.9.8.1. Minimum resource requirements for cluster installation	931
7.9.8.2. Tested instance types for Azure	933
7.9.8.3. Enabling trusted launch for Azure VMs	933
7.9.8.4. Enabling confidential VMs	934
7.9.8.5. Sample customized install-config.yaml file for Azure	935
7.9.8.6. Configuring the cluster-wide proxy during installation	939
7.9.9. Deploying the cluster	940
7.9.10. Installing the OpenShift CLI	943
Installing the OpenShift CLI on Linux	943
Installing the OpenShift CLI on Windows	944
Installing the OpenShift CLI on macOS	944
7.9.11. Logging in to the cluster by using the CLI	945
7.9.12. Telemetry access for OpenShift Container Platform	945
7.9.13. Next steps	946
7.10. INSTALLING A CLUSTER ON AZURE IN A RESTRICTED NETWORK WITH USER-PROVISIONED INFRASTRUCTURE	946
7.10.1. About installations in restricted networks	946
7.10.1.1. Additional limits	947
7.10.1.2. Internet access for OpenShift Container Platform	947
7.10.2. Configuring your Azure project	947
7.10.2.1. Azure account limits	948
7.10.2.2. Configuring a public DNS zone in Azure	951
7.10.2.3. Increasing Azure account limits	951
7.10.2.4. Certificate signing requests management	952
7.10.2.5. Required Azure roles	952
7.10.2.6. Required Azure permissions for user-provisioned infrastructure	953
7.10.2.7. Creating a service principal	959
7.10.2.8. Supported Azure regions	961
Supported Azure public regions	961
Supported Azure Government regions	963
7.10.3. Requirements for a cluster with user-provisioned infrastructure	963
7.10.3.1. Required machines for cluster installation	963
7.10.3.2. Minimum resource requirements for cluster installation	964
7.10.3.3. Tested instance types for Azure	965
7.10.3.4. Tested instance types for Azure on 64-bit ARM infrastructures	966
7.10.4. Using the Azure Marketplace offering	966
7.10.4.1. Obtaining the installation program	968

7.10.4.2. Generating a key pair for cluster node SSH access	969
7.10.5. Creating the installation files for Azure	971
7.10.5.1. Optional: Creating a separate /var partition	971
7.10.5.2. Creating the installation configuration file	973
7.10.5.3. Configuring the cluster-wide proxy during installation	977
7.10.5.4. Exporting common variables for ARM templates	979
7.10.5.5. Creating the Kubernetes manifest and Ignition config files	980
7.10.6. Creating the Azure resource group	983
7.10.7. Uploading the RHCOS cluster image and bootstrap Ignition config file	984
7.10.8. Example for creating DNS zones	985
7.10.9. Creating a VNet in Azure	986
7.10.9.1. ARM template for the VNet	987
7.10.10. Deploying the RHCOS cluster image for the Azure infrastructure	988
7.10.10.1. ARM template for image storage	989
7.10.11. Networking requirements for user-provisioned infrastructure	992
7.10.11.1. Network connectivity requirements	992
7.10.12. Creating networking and load balancing components in Azure	993
7.10.12.1. ARM template for the network and load balancers	995
7.10.13. Creating the bootstrap machine in Azure	999
7.10.13.1. ARM template for the bootstrap machine	1001
7.10.14. Creating the control plane machines in Azure	1004
7.10.14.1. ARM template for control plane machines	1005
7.10.15. Wait for bootstrap completion and remove bootstrap resources in Azure	1009
7.10.16. Creating additional worker machines in Azure	1010
7.10.16.1. ARM template for worker machines	1012
7.10.17. Installing the OpenShift CLI	1015
7.10.17.1. Installing the OpenShift CLI on Linux	1015
7.10.17.2. Installing the OpenShift CLI on Windows	1016
7.10.17.3. Installing the OpenShift CLI on macOS	1016
7.10.18. Logging in to the cluster by using the CLI	1017
7.10.19. Approving the certificate signing requests for your machines	1017
7.10.20. Adding the Ingress DNS records	1020
7.10.21. Completing an Azure installation on user-provisioned infrastructure	1021
7.10.22. Telemetry access for OpenShift Container Platform	1022
7.11. INSTALLING A CLUSTER ON AZURE USING ARM TEMPLATES	1022
7.11.1. Prerequisites	1023
7.11.2. Internet access for OpenShift Container Platform	1023
7.11.3. Configuring your Azure project	1024
7.11.3.1. Azure account limits	1024
7.11.3.2. Configuring a public DNS zone in Azure	1027
7.11.3.3. Increasing Azure account limits	1027
7.11.3.4. Certificate signing requests management	1028
7.11.3.5. Recording the subscription and tenant IDs	1028
7.11.3.6. Supported identities to access Azure resources	1030
7.11.3.7. Required Azure permissions for user-provisioned infrastructure	1030
7.11.3.8. Using Azure managed identities	1036
7.11.3.9. Creating a service principal	1037
7.11.3.10. Supported Azure regions	1038
7.11.3.10.1. Supported Azure public regions	1038
7.11.3.10.2. Supported Azure Government regions	1039
7.11.4. Requirements for a cluster with user-provisioned infrastructure	1040
7.11.4.1. Required machines for cluster installation	1040
7.11.4.2. Minimum resource requirements for cluster installation	1040

7.11.4.3. Tested instance types for Azure	1042
7.11.4.4. Tested instance types for Azure on 64-bit ARM infrastructures	1042
7.11.5. Using the Azure Marketplace offering	1042
7.11.6. Obtaining the installation program	1045
7.11.7. Generating a key pair for cluster node SSH access	1046
7.11.8. Creating the installation files for Azure	1048
7.11.8.1. Optional: Creating a separate /var partition	1048
7.11.8.2. Creating the installation configuration file	1050
7.11.8.3. Configuring the cluster-wide proxy during installation	1052
7.11.8.4. Exporting common variables for ARM templates	1054
7.11.8.5. Creating the Kubernetes manifest and Ignition config files	1055
7.11.9. Creating the Azure resource group	1058
7.11.10. Uploading the RHCOS cluster image and bootstrap Ignition config file	1059
7.11.11. Example for creating DNS zones	1061
7.11.12. Creating a VNet in Azure	1061
7.11.12.1. ARM template for the VNet	1062
7.11.13. Deploying the RHCOS cluster image for the Azure infrastructure	1064
7.11.13.1. ARM template for image storage	1065
7.11.14. Networking requirements for user-provisioned infrastructure	1068
7.11.14.1. Network connectivity requirements	1068
7.11.15. Creating networking and load balancing components in Azure	1069
7.11.15.1. ARM template for the network and load balancers	1070
7.11.16. Creating the bootstrap machine in Azure	1075
7.11.16.1. ARM template for the bootstrap machine	1076
7.11.17. Creating the control plane machines in Azure	1080
7.11.17.1. ARM template for control plane machines	1081
7.11.18. Wait for bootstrap completion and remove bootstrap resources in Azure	1085
7.11.19. Creating additional worker machines in Azure	1086
7.11.19.1. ARM template for worker machines	1087
7.11.20. Installing the OpenShift CLI	1091
Installing the OpenShift CLI on Linux	1091
Installing the OpenShift CLI on Windows	1091
Installing the OpenShift CLI on macOS	1092
7.11.21. Logging in to the cluster by using the CLI	1092
7.11.22. Approving the certificate signing requests for your machines	1093
7.11.23. Adding the Ingress DNS records	1096
7.11.24. Completing an Azure installation on user-provisioned infrastructure	1097
7.11.25. Telemetry access for OpenShift Container Platform	1098
7.12. INSTALLING A CLUSTER ON AZURE IN A RESTRICTED NETWORK	1098
7.12.1. Prerequisites	1098
7.12.2. About installations in restricted networks	1099
7.12.2.1. Additional limits	1099
7.12.2.2. User-defined outbound routing	1099
Restricted cluster with Azure Firewall	1100
7.12.3. About reusing a VNet for your OpenShift Container Platform cluster	1100
7.12.3.1. Requirements for using your VNet	1100
7.12.3.1.1. Network security group requirements	1102
7.12.3.2. Division of permissions	1103
7.12.3.3. Isolation between clusters	1103
7.12.4. Internet access for OpenShift Container Platform	1103
7.12.5. Generating a key pair for cluster node SSH access	1104
7.12.6. Creating the installation configuration file	1105
7.12.6.1. Minimum resource requirements for cluster installation	1109

7.12.6.2. Tested instance types for Azure	1110
7.12.6.3. Tested instance types for Azure on 64-bit ARM infrastructures	1111
7.12.6.4. Enabling trusted launch for Azure VMs	1111
7.12.6.5. Enabling confidential VMs	1112
7.12.6.6. Sample customized install-config.yaml file for Azure	1113
7.12.6.7. Configuring the cluster-wide proxy during installation	1117
7.12.7. Installing the OpenShift CLI	1118
Installing the OpenShift CLI on Linux	1118
Installing the OpenShift CLI on Windows	1119
Installing the OpenShift CLI on macOS	1119
7.12.8. Alternatives to storing administrator-level secrets in the kube-system project	1120
7.12.8.1. Manually creating long-term credentials	1120
7.12.8.2. Configuring an Azure cluster to use short-term credentials	1122
7.12.8.2.1. Configuring the Cloud Credential Operator utility	1122
7.12.8.2.2. Creating Azure resources with the Cloud Credential Operator utility	1125
7.12.8.2.3. Incorporating the Cloud Credential Operator utility manifests	1127
7.12.9. Deploying the cluster	1128
7.12.10. Logging in to the cluster by using the CLI	1129
7.12.11. Telemetry access for OpenShift Container Platform	1130
7.12.12. Next steps	1130
7.13. INSTALLING A THREE-NODE CLUSTER ON AZURE	1130
7.13.1. Configuring a three-node cluster	1131
7.13.2. Next steps	1132
7.14. UNINSTALLING A CLUSTER ON AZURE	1132
7.14.1. Removing a cluster that uses installer-provisioned infrastructure	1132
7.14.2. Deleting Microsoft Azure resources with the Cloud Credential Operator utility	1133
7.15. INSTALLATION CONFIGURATION PARAMETERS FOR AZURE	1133
7.15.1. Available installation configuration parameters for Azure	1133
7.15.1.1. Required configuration parameters	1134
7.15.1.2. Network configuration parameters	1135
7.15.1.3. Optional configuration parameters	1136
7.15.1.4. Additional Azure configuration parameters	1142
CHAPTER 8. INSTALLING ON AZURE STACK HUB	1161
8.1. PREPARING TO INSTALL ON AZURE STACK HUB	1161
8.1.1. Prerequisites	1161
8.1.2. Requirements for installing OpenShift Container Platform on Azure Stack Hub	1161
8.1.3. Choosing a method to install OpenShift Container Platform on Azure Stack Hub	1161
8.1.3.1. Installing a cluster on installer-provisioned infrastructure	1161
8.1.3.2. Installing a cluster on user-provisioned infrastructure	1161
8.1.4. Next steps	1161
8.2. CONFIGURING AN AZURE STACK HUB ACCOUNT	1162
8.2.1. Azure Stack Hub account limits	1162
8.2.2. Configuring a DNS zone in Azure Stack Hub	1163
8.2.3. Required Azure Stack Hub roles	1164
8.2.4. Creating a service principal	1164
8.2.5. Next steps	1166
8.3. INSTALLING A CLUSTER ON AZURE STACK HUB WITH AN INSTALLER-PROVISIONED INFRASTRUCTURE	1167
8.3.1. Prerequisites	1167
8.3.2. Internet access for OpenShift Container Platform	1167
8.3.3. Generating a key pair for cluster node SSH access	1168
8.3.4. Uploading the RHCOS cluster image	1169

8.3.5. Obtaining the installation program	1170
8.3.6. Manually creating the installation configuration file	1171
8.3.6.1. Sample customized install-config.yaml file for Azure Stack Hub	1172
8.3.7. Manually manage cloud credentials	1174
8.3.8. Configuring the cluster to use an internal CA	1176
8.3.9. Deploying the cluster	1176
8.3.10. Installing the OpenShift CLI	1178
Installing the OpenShift CLI on Linux	1178
Installing the OpenShift CLI on Windows	1179
Installing the OpenShift CLI on macOS	1179
8.3.11. Logging in to the cluster by using the CLI	1180
8.3.12. Logging in to the cluster by using the web console	1180
8.3.13. Telemetry access for OpenShift Container Platform	1181
8.3.14. Next steps	1181
8.4. INSTALLING A CLUSTER ON AZURE STACK HUB WITH NETWORK CUSTOMIZATIONS	1181
8.4.1. Prerequisites	1182
8.4.2. Internet access for OpenShift Container Platform	1182
8.4.3. Generating a key pair for cluster node SSH access	1182
8.4.4. Uploading the RHCOS cluster image	1184
8.4.5. Obtaining the installation program	1185
8.4.6. Manually creating the installation configuration file	1186
8.4.6.1. Sample customized install-config.yaml file for Azure Stack Hub	1187
8.4.7. Manually manage cloud credentials	1189
8.4.8. Configuring the cluster to use an internal CA	1191
8.4.9. Network configuration phases	1191
8.4.10. Specifying advanced network configuration	1192
8.4.11. Cluster Network Operator configuration	1193
8.4.11.1. Cluster Network Operator configuration object	1193
defaultNetwork object configuration	1194
Configuration for the OVN-Kubernetes network plugin	1195
kubeProxyConfig object configuration (OpenShiftSDN container network interface only)	1199
8.4.12. Configuring hybrid networking with OVN-Kubernetes	1200
8.4.13. Deploying the cluster	1201
8.4.14. Installing the OpenShift CLI	1202
Installing the OpenShift CLI on Linux	1203
Installing the OpenShift CLI on Windows	1203
Installing the OpenShift CLI on macOS	1204
8.4.15. Logging in to the cluster by using the CLI	1204
8.4.16. Logging in to the cluster by using the web console	1205
8.4.17. Telemetry access for OpenShift Container Platform	1206
8.4.18. Next steps	1206
8.5. INSTALLING A CLUSTER ON AZURE STACK HUB USING ARM TEMPLATES	1206
8.5.1. Prerequisites	1206
8.5.2. Internet access for OpenShift Container Platform	1207
8.5.3. Configuring your Azure Stack Hub project	1207
8.5.3.1. Azure Stack Hub account limits	1207
8.5.3.2. Configuring a DNS zone in Azure Stack Hub	1209
8.5.3.3. Certificate signing requests management	1209
8.5.3.4. Required Azure Stack Hub roles	1210
8.5.3.5. Creating a service principal	1210
8.5.4. Obtaining the installation program	1212
8.5.5. Generating a key pair for cluster node SSH access	1213
8.5.6. Creating the installation files for Azure Stack Hub	1215

8.5.6.1. Manually creating the installation configuration file	1215
8.5.6.2. Sample customized install-config.yaml file for Azure Stack Hub	1217
8.5.6.3. Configuring the cluster-wide proxy during installation	1219
8.5.6.4. Exporting common variables for ARM templates	1220
8.5.6.5. Creating the Kubernetes manifest and Ignition config files	1222
8.5.6.6. Optional: Creating a separate /var partition	1226
8.5.7. Creating the Azure resource group	1228
8.5.8. Uploading the RHCOS cluster image and bootstrap Ignition config file	1229
8.5.9. Example for creating DNS zones	1230
8.5.10. Creating a VNet in Azure Stack Hub	1231
8.5.10.1. ARM template for the VNet	1232
8.5.11. Deploying the RHCOS cluster image for the Azure Stack Hub infrastructure	1232
8.5.11.1. ARM template for image storage	1233
8.5.12. Networking requirements for user-provisioned infrastructure	1233
8.5.12.1. Network connectivity requirements	1233
8.5.13. Creating networking and load balancing components in Azure Stack Hub	1234
8.5.13.1. ARM template for the network and load balancers	1236
8.5.14. Creating the bootstrap machine in Azure Stack Hub	1236
8.5.14.1. ARM template for the bootstrap machine	1237
8.5.15. Creating the control plane machines in Azure Stack Hub	1238
8.5.15.1. ARM template for control plane machines	1238
8.5.16. Wait for bootstrap completion and remove bootstrap resources in Azure Stack Hub	1239
8.5.17. Creating additional worker machines in Azure Stack Hub	1240
8.5.17.1. ARM template for worker machines	1241
8.5.18. Installing the OpenShift CLI	1241
Installing the OpenShift CLI on Linux	1241
Installing the OpenShift CLI on Windows	1242
Installing the OpenShift CLI on macOS	1242
8.5.19. Logging in to the cluster by using the CLI	1243
8.5.20. Approving the certificate signing requests for your machines	1243
8.5.21. Adding the Ingress DNS records	1246
8.5.22. Completing an Azure Stack Hub installation on user-provisioned infrastructure	1247
8.6. INSTALLATION CONFIGURATION PARAMETERS FOR AZURE STACK HUB	1248
8.6.1. Available installation configuration parameters for Azure Stack Hub	1248
8.6.1.1. Required configuration parameters	1248
8.6.1.2. Network configuration parameters	1250
8.6.1.3. Optional configuration parameters	1251
8.6.1.4. Additional Azure Stack Hub configuration parameters	1257
8.7. UNINSTALLING A CLUSTER ON AZURE STACK HUB	1260
8.7.1. Removing a cluster that uses installer-provisioned infrastructure	1260
CHAPTER 9. INSTALLING ON GCP	1262
9.1. PREPARING TO INSTALL ON GCP	1262
9.1.1. Prerequisites	1262
9.1.2. Requirements for installing OpenShift Container Platform on GCP	1262
9.1.3. Choosing a method to install OpenShift Container Platform on GCP	1262
9.1.3.1. Installing a cluster on installer-provisioned infrastructure	1262
9.1.3.2. Installing a cluster on user-provisioned infrastructure	1263
9.1.4. Next steps	1263
9.2. CONFIGURING A GCP PROJECT	1263
9.2.1. Creating a GCP project	1263
9.2.2. Enabling API services in GCP	1264
9.2.3. Configuring DNS for GCP	1265

9.2.4. GCP account limits	1265
9.2.5. Creating a service account in GCP	1267
9.2.5.1. Required GCP roles	1268
9.2.5.2. Required GCP permissions for installer-provisioned infrastructure	1269
9.2.5.3. Required GCP permissions for shared VPC installations	1275
9.2.6. Supported GCP regions	1276
9.2.7. Next steps	1278
9.3. INSTALLING A CLUSTER QUICKLY ON GCP	1278
9.3.1. Prerequisites	1278
9.3.2. Internet access for OpenShift Container Platform	1278
9.3.3. Generating a key pair for cluster node SSH access	1278
9.3.4. Obtaining the installation program	1280
9.3.5. Deploying the cluster	1281
9.3.6. Installing the OpenShift CLI	1284
Installing the OpenShift CLI on Linux	1284
Installing the OpenShift CLI on Windows	1284
Installing the OpenShift CLI on macOS	1285
9.3.7. Logging in to the cluster by using the CLI	1285
9.3.8. Telemetry access for OpenShift Container Platform	1286
9.3.9. Next steps	1286
9.4. INSTALLING A CLUSTER ON GCP WITH CUSTOMIZATIONS	1286
9.4.1. Prerequisites	1286
9.4.2. Internet access for OpenShift Container Platform	1287
9.4.3. Generating a key pair for cluster node SSH access	1287
9.4.4. Obtaining the installation program	1289
9.4.5. Creating the installation configuration file	1290
9.4.5.1. Minimum resource requirements for cluster installation	1291
9.4.5.2. Tested instance types for GCP	1292
9.4.5.3. Tested instance types for GCP on 64-bit ARM infrastructures	1293
9.4.5.4. Using custom machine types	1293
9.4.5.5. Enabling Shielded VMs	1294
9.4.5.6. Enabling Confidential VMs	1294
9.4.5.7. Sample customized install-config.yaml file for GCP	1295
9.4.5.8. Configuring the cluster-wide proxy during installation	1298
9.4.6. Managing user-defined labels and tags for GCP	1300
9.4.6.1. Configuring user-defined labels and tags for GCP	1302
9.4.6.2. Querying user-defined labels and tags for GCP	1303
9.4.7. Installing the OpenShift CLI	1304
Installing the OpenShift CLI on Linux	1304
Installing the OpenShift CLI on Windows	1305
Installing the OpenShift CLI on macOS	1305
9.4.8. Alternatives to storing administrator-level secrets in the kube-system project	1306
9.4.8.1. Manually creating long-term credentials	1306
9.4.8.2. Configuring a GCP cluster to use short-term credentials	1308
9.4.8.2.1. Configuring the Cloud Credential Operator utility	1309
9.4.8.2.2. Creating GCP resources with the Cloud Credential Operator utility	1311
9.4.8.2.3. Incorporating the Cloud Credential Operator utility manifests	1313
9.4.9. Using the GCP Marketplace offering	1314
9.4.10. Deploying the cluster	1315
9.4.11. Logging in to the cluster by using the CLI	1317
9.4.12. Telemetry access for OpenShift Container Platform	1317
9.4.13. Next steps	1318
9.5. INSTALLING A CLUSTER ON GCP WITH NETWORK CUSTOMIZATIONS	1318

9.5.1. Prerequisites	1318
9.5.2. Internet access for OpenShift Container Platform	1318
9.5.3. Generating a key pair for cluster node SSH access	1319
9.5.4. Obtaining the installation program	1321
9.5.5. Creating the installation configuration file	1321
9.5.5.1. Minimum resource requirements for cluster installation	1323
9.5.5.2. Tested instance types for GCP	1324
9.5.5.3. Tested instance types for GCP on 64-bit ARM infrastructures	1324
9.5.5.4. Using custom machine types	1325
9.5.5.5. Enabling Shielded VMs	1325
9.5.5.6. Enabling Confidential VMs	1326
9.5.5.7. Sample customized install-config.yaml file for GCP	1327
9.5.5.8. Configuring the cluster-wide proxy during installation	1330
9.5.6. Installing the OpenShift CLI	1332
Installing the OpenShift CLI on Linux	1332
Installing the OpenShift CLI on Windows	1333
Installing the OpenShift CLI on macOS	1333
9.5.7. Alternatives to storing administrator-level secrets in the kube-system project	1334
9.5.7.1. Manually creating long-term credentials	1334
9.5.7.2. Configuring a GCP cluster to use short-term credentials	1336
9.5.7.2.1. Configuring the Cloud Credential Operator utility	1336
9.5.7.2.2. Creating GCP resources with the Cloud Credential Operator utility	1339
9.5.7.2.3. Incorporating the Cloud Credential Operator utility manifests	1341
9.5.8. Network configuration phases	1342
9.5.9. Specifying advanced network configuration	1343
9.5.10. Cluster Network Operator configuration	1344
9.5.10.1. Cluster Network Operator configuration object	1344
defaultNetwork object configuration	1345
Configuration for the OVN-Kubernetes network plugin	1346
kubeProxyConfig object configuration (OpenShiftSDN container network interface only)	1350
9.5.11. Deploying the cluster	1351
9.5.12. Logging in to the cluster by using the CLI	1353
9.5.13. Telemetry access for OpenShift Container Platform	1353
9.5.14. Next steps	1353
9.6. INSTALLING A CLUSTER ON GCP IN A RESTRICTED NETWORK	1354
9.6.1. Prerequisites	1354
9.6.2. About installations in restricted networks	1354
9.6.2.1. Additional limits	1355
9.6.3. Internet access for OpenShift Container Platform	1355
9.6.4. Generating a key pair for cluster node SSH access	1355
9.6.5. Creating the installation configuration file	1357
9.6.5.1. Minimum resource requirements for cluster installation	1359
9.6.5.2. Tested instance types for GCP	1360
9.6.5.3. Tested instance types for GCP on 64-bit ARM infrastructures	1361
9.6.5.4. Using custom machine types	1361
9.6.5.5. Enabling Shielded VMs	1362
9.6.5.6. Enabling Confidential VMs	1363
9.6.5.7. Sample customized install-config.yaml file for GCP	1364
9.6.5.8. Create an Ingress Controller with global access on GCP	1367
9.6.5.9. Configuring the cluster-wide proxy during installation	1368
9.6.6. Installing the OpenShift CLI	1370
Installing the OpenShift CLI on Linux	1370
Installing the OpenShift CLI on Windows	1371

Installing the OpenShift CLI on macOS	1371
9.6.7. Alternatives to storing administrator-level secrets in the kube-system project	1372
9.6.7.1. Manually creating long-term credentials	1372
9.6.7.2. Configuring a GCP cluster to use short-term credentials	1374
9.6.7.2.1. Configuring the Cloud Credential Operator utility	1374
9.6.7.2.2. Creating GCP resources with the Cloud Credential Operator utility	1377
9.6.7.2.3. Incorporating the Cloud Credential Operator utility manifests	1379
9.6.8. Deploying the cluster	1380
9.6.9. Logging in to the cluster by using the CLI	1382
9.6.10. Disabling the default OperatorHub catalog sources	1383
9.6.11. Telemetry access for OpenShift Container Platform	1383
9.6.12. Next steps	1383
9.7. INSTALLING A CLUSTER ON GCP INTO AN EXISTING VPC	1383
9.7.1. Prerequisites	1384
9.7.2. About using a custom VPC	1384
9.7.2.1. Requirements for using your VPC	1384
9.7.2.2. VPC validation	1384
9.7.2.3. Division of permissions	1385
9.7.2.4. Isolation between clusters	1385
9.7.3. Internet access for OpenShift Container Platform	1385
9.7.4. Generating a key pair for cluster node SSH access	1385
9.7.5. Obtaining the installation program	1387
9.7.6. Creating the installation configuration file	1388
9.7.6.1. Minimum resource requirements for cluster installation	1389
9.7.6.2. Tested instance types for GCP	1390
9.7.6.3. Tested instance types for GCP on 64-bit ARM infrastructures	1391
9.7.6.4. Using custom machine types	1391
9.7.6.5. Enabling Shielded VMs	1392
9.7.6.6. Enabling Confidential VMs	1393
9.7.6.7. Sample customized install-config.yaml file for GCP	1394
9.7.6.8. Create an Ingress Controller with global access on GCP	1397
9.7.6.9. Configuring the cluster-wide proxy during installation	1398
9.7.7. Installing the OpenShift CLI	1400
Installing the OpenShift CLI on Linux	1400
Installing the OpenShift CLI on Windows	1401
Installing the OpenShift CLI on macOS	1401
9.7.8. Alternatives to storing administrator-level secrets in the kube-system project	1402
9.7.8.1. Manually creating long-term credentials	1402
9.7.8.2. Configuring a GCP cluster to use short-term credentials	1404
9.7.8.2.1. Configuring the Cloud Credential Operator utility	1404
9.7.8.2.2. Creating GCP resources with the Cloud Credential Operator utility	1407
9.7.8.2.3. Incorporating the Cloud Credential Operator utility manifests	1409
9.7.9. Deploying the cluster	1410
9.7.10. Logging in to the cluster by using the CLI	1412
9.7.11. Telemetry access for OpenShift Container Platform	1412
9.7.12. Next steps	1413
9.8. INSTALLING A CLUSTER ON GCP INTO A SHARED VPC	1413
9.8.1. Prerequisites	1413
9.8.2. Internet access for OpenShift Container Platform	1413
9.8.3. Generating a key pair for cluster node SSH access	1414
9.8.4. Obtaining the installation program	1416
9.8.5. Creating the installation files for GCP	1416
9.8.5.1. Manually creating the installation configuration file	1417

9.8.5.2. Enabling Shielded VMs	1417
9.8.5.3. Enabling Confidential VMs	1418
9.8.5.4. Sample customized install-config.yaml file for shared VPC installation	1419
9.8.5.5. Configuring the cluster-wide proxy during installation	1421
9.8.6. Installing the OpenShift CLI	1422
Installing the OpenShift CLI on Linux	1422
Installing the OpenShift CLI on Windows	1423
Installing the OpenShift CLI on macOS	1423
9.8.7. Alternatives to storing administrator-level secrets in the kube-system project	1424
9.8.7.1. Manually creating long-term credentials	1424
9.8.7.2. Configuring a GCP cluster to use short-term credentials	1427
9.8.7.2.1. Configuring the Cloud Credential Operator utility	1427
9.8.7.2.2. Creating GCP resources with the Cloud Credential Operator utility	1429
9.8.7.2.3. Incorporating the Cloud Credential Operator utility manifests	1431
9.8.8. Deploying the cluster	1432
9.8.9. Logging in to the cluster by using the CLI	1434
9.8.10. Telemetry access for OpenShift Container Platform	1435
9.8.11. Next steps	1435
9.9. INSTALLING A PRIVATE CLUSTER ON GCP	1435
9.9.1. Prerequisites	1435
9.9.2. Private clusters	1435
9.9.2.1. Private clusters in GCP	1436
9.9.2.1.1. Limitations	1437
9.9.3. About using a custom VPC	1437
9.9.3.1. Requirements for using your VPC	1437
9.9.3.2. Division of permissions	1438
9.9.3.3. Isolation between clusters	1438
9.9.4. Internet access for OpenShift Container Platform	1438
9.9.5. Generating a key pair for cluster node SSH access	1439
9.9.6. Obtaining the installation program	1441
9.9.7. Manually creating the installation configuration file	1441
9.9.7.1. Minimum resource requirements for cluster installation	1442
9.9.7.2. Tested instance types for GCP	1443
9.9.7.3. Tested instance types for GCP on 64-bit ARM infrastructures	1444
9.9.7.4. Using custom machine types	1444
9.9.7.5. Enabling Shielded VMs	1445
9.9.7.6. Enabling Confidential VMs	1446
9.9.7.7. Sample customized install-config.yaml file for GCP	1447
9.9.7.8. Create an Ingress Controller with global access on GCP	1450
9.9.7.9. Configuring the cluster-wide proxy during installation	1451
9.9.8. Installing the OpenShift CLI	1453
Installing the OpenShift CLI on Linux	1453
Installing the OpenShift CLI on Windows	1454
Installing the OpenShift CLI on macOS	1454
9.9.9. Alternatives to storing administrator-level secrets in the kube-system project	1455
9.9.9.1. Manually creating long-term credentials	1455
9.9.9.2. Configuring a GCP cluster to use short-term credentials	1457
9.9.9.2.1. Configuring the Cloud Credential Operator utility	1457
9.9.9.2.2. Creating GCP resources with the Cloud Credential Operator utility	1460
9.9.9.2.3. Incorporating the Cloud Credential Operator utility manifests	1462
9.9.10. Deploying the cluster	1463
9.9.11. Logging in to the cluster by using the CLI	1465
9.9.12. Telemetry access for OpenShift Container Platform	1465

9.9.13. Next steps	1466
9.10. INSTALLING A CLUSTER ON USER-PROVISIONED INFRASTRUCTURE IN GCP BY USING DEPLOYMENT MANAGER TEMPLATES	1466
9.10.1. Prerequisites	1466
9.10.2. Certificate signing requests management	1467
9.10.3. Internet access for OpenShift Container Platform	1467
9.10.4. Configuring your GCP project	1467
9.10.4.1. Creating a GCP project	1467
9.10.4.2. Enabling API services in GCP	1468
9.10.4.3. Configuring DNS for GCP	1469
9.10.4.4. GCP account limits	1469
9.10.4.5. Creating a service account in GCP	1471
9.10.4.6. Required GCP roles	1471
9.10.4.7. Required GCP permissions for user-provisioned infrastructure	1473
9.10.4.8. Supported GCP regions	1480
9.10.4.9. Installing and configuring CLI tools for GCP	1481
9.10.5. Requirements for a cluster with user-provisioned infrastructure	1482
9.10.5.1. Required machines for cluster installation	1482
9.10.5.2. Minimum resource requirements for cluster installation	1483
9.10.5.3. Tested instance types for GCP	1484
9.10.5.4. Tested instance types for GCP on 64-bit ARM infrastructures	1484
9.10.5.5. Using custom machine types	1485
9.10.6. Creating the installation files for GCP	1485
9.10.6.1. Optional: Creating a separate /var partition	1485
9.10.6.2. Creating the installation configuration file	1487
9.10.6.3. Enabling Shielded VMs	1489
9.10.6.4. Enabling Confidential VMs	1490
9.10.6.5. Configuring the cluster-wide proxy during installation	1491
9.10.6.6. Creating the Kubernetes manifest and Ignition config files	1492
9.10.7. Exporting common variables	1495
9.10.7.1. Extracting the infrastructure name	1495
9.10.7.2. Exporting common variables for Deployment Manager templates	1496
9.10.8. Creating a VPC in GCP	1496
9.10.8.1. Deployment Manager template for the VPC	1497
9.10.9. Networking requirements for user-provisioned infrastructure	1499
9.10.9.1. Setting the cluster node hostnames through DHCP	1499
9.10.9.2. Network connectivity requirements	1499
9.10.10. Creating load balancers in GCP	1500
9.10.10.1. Deployment Manager template for the external load balancer	1502
9.10.10.2. Deployment Manager template for the internal load balancer	1503
9.10.11. Creating a private DNS zone in GCP	1504
9.10.11.1. Deployment Manager template for the private DNS	1506
9.10.12. Creating firewall rules in GCP	1506
9.10.12.1. Deployment Manager template for firewall rules	1507
9.10.13. Creating IAM roles in GCP	1510
9.10.13.1. Deployment Manager template for IAM roles	1512
9.10.14. Creating the RHCOS cluster image for the GCP infrastructure	1512
9.10.15. Creating the bootstrap machine in GCP	1513
9.10.15.1. Deployment Manager template for the bootstrap machine	1515
9.10.16. Creating the control plane machines in GCP	1516
9.10.16.1. Deployment Manager template for control plane machines	1518
9.10.17. Wait for bootstrap completion and remove bootstrap resources in GCP	1520
9.10.18. Creating additional worker machines in GCP	1521

9.10.18.1. Deployment Manager template for worker machines	1524
9.10.19. Installing the OpenShift CLI	1525
Installing the OpenShift CLI on Linux	1525
Installing the OpenShift CLI on Windows	1525
Installing the OpenShift CLI on macOS	1526
9.10.20. Logging in to the cluster by using the CLI	1526
9.10.21. Approving the certificate signing requests for your machines	1527
9.10.22. Optional: Adding the ingress DNS records	1530
9.10.23. Completing a GCP installation on user-provisioned infrastructure	1531
9.10.24. Telemetry access for OpenShift Container Platform	1534
9.10.25. Next steps	1534
9.11. INSTALLING A CLUSTER INTO A SHARED VPC ON GCP USING DEPLOYMENT MANAGER TEMPLATES	1534
9.11.1. Prerequisites	1534
9.11.2. Certificate signing requests management	1535
9.11.3. Internet access for OpenShift Container Platform	1535
9.11.4. Configuring the GCP project that hosts your cluster	1535
9.11.4.1. Creating a GCP project	1535
9.11.4.2. Enabling API services in GCP	1536
9.11.4.3. GCP account limits	1537
9.11.4.4. Creating a service account in GCP	1538
9.11.4.4.1. Required GCP roles	1539
9.11.4.5. Supported GCP regions	1540
9.11.4.6. Installing and configuring CLI tools for GCP	1541
9.11.5. Requirements for a cluster with user-provisioned infrastructure	1542
9.11.5.1. Required machines for cluster installation	1542
9.11.5.2. Minimum resource requirements for cluster installation	1543
9.11.5.3. Tested instance types for GCP	1544
9.11.5.4. Using custom machine types	1544
9.11.6. Configuring the GCP project that hosts your shared VPC network	1545
9.11.6.1. Configuring DNS for GCP	1546
9.11.6.2. Creating a VPC in GCP	1547
9.11.6.2.1. Deployment Manager template for the VPC	1548
9.11.7. Creating the installation files for GCP	1549
9.11.7.1. Manually creating the installation configuration file	1550
9.11.7.2. Enabling Shielded VMs	1551
9.11.7.3. Enabling Confidential VMs	1551
9.11.7.4. Sample customized install-config.yaml file for GCP	1552
9.11.7.5. Configuring the cluster-wide proxy during installation	1555
9.11.7.6. Creating the Kubernetes manifest and Ignition config files	1556
9.11.8. Exporting common variables	1559
9.11.8.1. Extracting the infrastructure name	1559
9.11.8.2. Exporting common variables for Deployment Manager templates	1560
9.11.9. Networking requirements for user-provisioned infrastructure	1561
9.11.9.1. Setting the cluster node hostnames through DHCP	1561
9.11.9.2. Network connectivity requirements	1561
9.11.10. Creating load balancers in GCP	1562
9.11.10.1. Deployment Manager template for the external load balancer	1564
9.11.10.2. Deployment Manager template for the internal load balancer	1565
9.11.11. Creating a private DNS zone in GCP	1567
9.11.11.1. Deployment Manager template for the private DNS	1568
9.11.12. Creating firewall rules in GCP	1569
9.11.12.1. Deployment Manager template for firewall rules	1570

9.11.13. Creating IAM roles in GCP	1572
9.11.13.1. Deployment Manager template for IAM roles	1575
9.11.14. Creating the RHCOS cluster image for the GCP infrastructure	1575
9.11.15. Creating the bootstrap machine in GCP	1576
9.11.15.1. Deployment Manager template for the bootstrap machine	1578
9.11.16. Creating the control plane machines in GCP	1579
9.11.16.1. Deployment Manager template for control plane machines	1581
9.11.17. Wait for bootstrap completion and remove bootstrap resources in GCP	1583
9.11.18. Creating additional worker machines in GCP	1584
9.11.18.1. Deployment Manager template for worker machines	1586
9.11.19. Installing the OpenShift CLI	1587
Installing the OpenShift CLI on Linux	1588
Installing the OpenShift CLI on Windows	1588
Installing the OpenShift CLI on macOS	1589
9.11.20. Logging in to the cluster by using the CLI	1589
9.11.21. Approving the certificate signing requests for your machines	1590
9.11.22. Adding the ingress DNS records	1592
9.11.23. Adding ingress firewall rules	1594
9.11.23.1. Creating cluster-wide firewall rules for a shared VPC in GCP	1595
9.11.24. Completing a GCP installation on user-provisioned infrastructure	1595
9.11.25. Telemetry access for OpenShift Container Platform	1598
9.11.26. Next steps	1598
9.12. INSTALLING A CLUSTER ON GCP IN A RESTRICTED NETWORK WITH USER-PROVISIONED INFRASTRUCTURE	1598
9.12.1. Prerequisites	1599
9.12.2. About installations in restricted networks	1599
9.12.2.1. Additional limits	1600
9.12.3. Internet access for OpenShift Container Platform	1600
9.12.4. Configuring your GCP project	1600
9.12.4.1. Creating a GCP project	1600
9.12.4.2. Enabling API services in GCP	1601
9.12.4.3. Configuring DNS for GCP	1601
9.12.4.4. GCP account limits	1602
9.12.4.5. Creating a service account in GCP	1604
9.12.4.6. Required GCP roles	1604
9.12.4.7. Required GCP permissions for user-provisioned infrastructure	1605
9.12.4.8. Supported GCP regions	1613
9.12.4.9. Installing and configuring CLI tools for GCP	1614
9.12.5. Requirements for a cluster with user-provisioned infrastructure	1615
9.12.5.1. Required machines for cluster installation	1615
9.12.5.2. Minimum resource requirements for cluster installation	1615
9.12.5.3. Tested instance types for GCP	1616
9.12.5.4. Using custom machine types	1617
9.12.6. Creating the installation files for GCP	1617
9.12.6.1. Optional: Creating a separate /var partition	1617
9.12.6.2. Creating the installation configuration file	1620
9.12.6.3. Enabling Shielded VMs	1622
9.12.6.4. Enabling Confidential VMs	1623
9.12.6.5. Configuring the cluster-wide proxy during installation	1624
9.12.6.6. Creating the Kubernetes manifest and Ignition config files	1626
9.12.7. Exporting common variables	1628
9.12.7.1. Extracting the infrastructure name	1628
9.12.7.2. Exporting common variables for Deployment Manager templates	1629

9.12.8. Creating a VPC in GCP	1629
9.12.8.1. Deployment Manager template for the VPC	1630
9.12.9. Networking requirements for user-provisioned infrastructure	1632
9.12.9.1. Setting the cluster node hostnames through DHCP	1632
9.12.9.2. Network connectivity requirements	1632
9.12.10. Creating load balancers in GCP	1633
9.12.10.1. Deployment Manager template for the external load balancer	1635
9.12.10.2. Deployment Manager template for the internal load balancer	1636
9.12.11. Creating a private DNS zone in GCP	1637
9.12.11.1. Deployment Manager template for the private DNS	1639
9.12.12. Creating firewall rules in GCP	1639
9.12.12.1. Deployment Manager template for firewall rules	1640
9.12.13. Creating IAM roles in GCP	1643
9.12.13.1. Deployment Manager template for IAM roles	1645
9.12.14. Creating the RHCOS cluster image for the GCP infrastructure	1645
9.12.15. Creating the bootstrap machine in GCP	1646
9.12.15.1. Deployment Manager template for the bootstrap machine	1648
9.12.16. Creating the control plane machines in GCP	1649
9.12.16.1. Deployment Manager template for control plane machines	1651
9.12.17. Wait for bootstrap completion and remove bootstrap resources in GCP	1653
9.12.18. Creating additional worker machines in GCP	1654
9.12.18.1. Deployment Manager template for worker machines	1656
9.12.19. Logging in to the cluster by using the CLI	1657
9.12.20. Disabling the default OperatorHub catalog sources	1658
9.12.21. Approving the certificate signing requests for your machines	1658
9.12.22. Optional: Adding the ingress DNS records	1661
9.12.23. Completing a GCP installation on user-provisioned infrastructure	1663
9.12.24. Telemetry access for OpenShift Container Platform	1665
9.12.25. Next steps	1665
9.13. INSTALLING A THREE-NODE CLUSTER ON GCP	1665
9.13.1. Configuring a three-node cluster	1666
9.13.2. Next steps	1667
9.14. INSTALLATION CONFIGURATION PARAMETERS FOR GCP	1667
9.14.1. Available installation configuration parameters for GCP	1667
9.14.1.1. Required configuration parameters	1667
9.14.1.2. Network configuration parameters	1669
9.14.1.3. Optional configuration parameters	1670
9.14.1.4. Additional Google Cloud Platform (GCP) configuration parameters	1677
9.15. UNINSTALLING A CLUSTER ON GCP	1709
9.15.1. Removing a cluster that uses installer-provisioned infrastructure	1709
9.15.2. Deleting Google Cloud Platform resources with the Cloud Credential Operator utility	1710
CHAPTER 10. INSTALLING ON IBM CLOUD	1712
10.1. PREPARING TO INSTALL ON IBM CLOUD	1712
10.1.1. Prerequisites	1712
10.1.2. Requirements for installing OpenShift Container Platform on IBM Cloud	1712
10.1.3. Choosing a method to install OpenShift Container Platform on IBM Cloud	1712
10.1.3.1. Installing a cluster on installer-provisioned infrastructure	1712
10.1.4. Next steps	1713
10.2. CONFIGURING AN IBM CLOUD ACCOUNT	1713
10.2.1. Prerequisites	1713
10.2.2. Quotas and limits on IBM Cloud	1713
Virtual Private Cloud (VPC)	1713

Application load balancer	1713
Floating IP address	1713
Virtual Server Instances (VSI)	1714
Block Storage Volumes	1714
10.2.3. Configuring DNS resolution	1715
10.2.3.1. Using IBM Cloud Internet Services for DNS resolution	1715
10.2.3.2. Using IBM Cloud DNS Services for DNS resolution	1716
10.2.4. IBM Cloud IAM Policies and API Key	1717
10.2.4.1. Required access policies	1717
10.2.4.2. Access policy assignment	1718
10.2.4.3. Creating an API key	1718
10.2.5. Supported IBM Cloud regions	1719
10.2.6. Next steps	1719
10.3. CONFIGURING IAM FOR IBM CLOUD	1719
10.3.1. Alternatives to storing administrator-level secrets in the kube-system project	1719
10.3.2. Configuring the Cloud Credential Operator utility	1720
10.3.3. Next steps	1721
10.3.4. Additional resources	1721
10.4. USER-MANAGED ENCRYPTION FOR IBM CLOUD	1721
10.4.1. Next steps	1722
10.5. INSTALLING A CLUSTER ON IBM CLOUD WITH CUSTOMIZATIONS	1722
10.5.1. Prerequisites	1722
10.5.2. Internet access for OpenShift Container Platform	1723
10.5.3. Generating a key pair for cluster node SSH access	1723
10.5.4. Obtaining the installation program	1725
10.5.5. Exporting the API key	1726
10.5.6. Creating the installation configuration file	1726
10.5.6.1. Minimum resource requirements for cluster installation	1727
10.5.6.2. Tested instance types for IBM Cloud	1728
10.5.6.3. Sample customized install-config.yaml file for IBM Cloud	1729
10.5.6.4. Configuring the cluster-wide proxy during installation	1731
10.5.7. Manually creating IAM	1732
10.5.8. Deploying the cluster	1735
10.5.9. Installing the OpenShift CLI	1736
Installing the OpenShift CLI on Linux	1736
Installing the OpenShift CLI on Windows	1737
Installing the OpenShift CLI on macOS	1737
10.5.10. Logging in to the cluster by using the CLI	1738
10.5.11. Telemetry access for OpenShift Container Platform	1739
10.5.12. Next steps	1739
10.6. INSTALLING A CLUSTER ON IBM CLOUD WITH NETWORK CUSTOMIZATIONS	1739
10.6.1. Prerequisites	1739
10.6.2. Internet access for OpenShift Container Platform	1740
10.6.3. Generating a key pair for cluster node SSH access	1740
10.6.4. Obtaining the installation program	1742
10.6.5. Exporting the API key	1743
10.6.6. Creating the installation configuration file	1743
10.6.6.1. Minimum resource requirements for cluster installation	1744
10.6.6.2. Tested instance types for IBM Cloud	1745
10.6.6.3. Sample customized install-config.yaml file for IBM Cloud	1746
10.6.6.4. Configuring the cluster-wide proxy during installation	1747
10.6.7. Manually creating IAM	1749
10.6.8. Network configuration phases	1751

10.6.9. Specifying advanced network configuration	1752
10.6.10. Cluster Network Operator configuration	1753
10.6.10.1. Cluster Network Operator configuration object	1754
defaultNetwork object configuration	1754
Configuration for the OVN-Kubernetes network plugin	1755
kubeProxyConfig object configuration (OpenShiftSDN container network interface only)	1759
10.6.11. Deploying the cluster	1760
10.6.12. Installing the OpenShift CLI	1761
Installing the OpenShift CLI on Linux	1761
Installing the OpenShift CLI on Windows	1762
Installing the OpenShift CLI on macOS	1762
10.6.13. Logging in to the cluster by using the CLI	1763
10.6.14. Telemetry access for OpenShift Container Platform	1764
10.6.15. Next steps	1764
10.7. INSTALLING A CLUSTER ON IBM CLOUD INTO AN EXISTING VPC	1764
10.7.1. Prerequisites	1764
10.7.2. About using a custom VPC	1764
10.7.2.1. Requirements for using your VPC	1765
10.7.2.2. VPC validation	1765
10.7.2.3. Isolation between clusters	1766
10.7.3. Internet access for OpenShift Container Platform	1766
10.7.4. Generating a key pair for cluster node SSH access	1766
10.7.5. Obtaining the installation program	1768
10.7.6. Exporting the API key	1769
10.7.7. Creating the installation configuration file	1770
10.7.7.1. Minimum resource requirements for cluster installation	1771
10.7.7.2. Tested instance types for IBM Cloud	1771
10.7.7.3. Sample customized install-config.yaml file for IBM Cloud	1772
10.7.7.4. Configuring the cluster-wide proxy during installation	1774
10.7.8. Manually creating IAM	1776
10.7.9. Deploying the cluster	1779
10.7.10. Installing the OpenShift CLI	1780
Installing the OpenShift CLI on Linux	1780
Installing the OpenShift CLI on Windows	1781
Installing the OpenShift CLI on macOS	1781
10.7.11. Logging in to the cluster by using the CLI	1782
10.7.12. Telemetry access for OpenShift Container Platform	1782
10.7.13. Next steps	1783
10.8. INSTALLING A PRIVATE CLUSTER ON IBM CLOUD	1783
10.8.1. Prerequisites	1783
10.8.2. Private clusters	1783
10.8.3. Private clusters in IBM Cloud	1784
10.8.3.1. Limitations	1784
10.8.4. About using a custom VPC	1784
10.8.4.1. Requirements for using your VPC	1784
10.8.4.2. VPC validation	1785
10.8.4.3. Isolation between clusters	1786
10.8.5. Internet access for OpenShift Container Platform	1786
10.8.6. Generating a key pair for cluster node SSH access	1786
10.8.7. Obtaining the installation program	1788
10.8.8. Exporting the API key	1789
10.8.9. Manually creating the installation configuration file	1789
10.8.9.1. Minimum resource requirements for cluster installation	1790

10.8.9.2. Tested instance types for IBM Cloud	1791
10.8.9.3. Sample customized install-config.yaml file for IBM Cloud	1792
10.8.9.4. Configuring the cluster-wide proxy during installation	1794
10.8.10. Manually creating IAM	1796
10.8.11. Deploying the cluster	1798
10.8.12. Installing the OpenShift CLI	1800
Installing the OpenShift CLI on Linux	1800
Installing the OpenShift CLI on Windows	1801
Installing the OpenShift CLI on macOS	1801
10.8.13. Logging in to the cluster by using the CLI	1802
10.8.14. Telemetry access for OpenShift Container Platform	1802
10.8.15. Next steps	1802
10.9. INSTALLING A CLUSTER ON IBM CLOUD IN A RESTRICTED NETWORK	1803
10.9.1. Prerequisites	1803
10.9.2. About installations in restricted networks	1803
10.9.2.1. Required internet access and an installation host	1804
10.9.2.2. Access to a mirror registry	1804
10.9.2.3. Access to IBM service endpoints	1804
10.9.2.4. Additional limits	1805
10.9.3. About using a custom VPC	1805
10.9.3.1. Requirements for using your VPC	1805
10.9.3.2. VPC validation	1805
10.9.3.3. Isolation between clusters	1806
10.9.3.4. Allowing endpoint gateway traffic	1806
10.9.4. Generating a key pair for cluster node SSH access	1807
10.9.5. Exporting the API key	1809
10.9.6. Downloading the RHCOS cluster image	1809
10.9.7. Manually creating the installation configuration file	1810
10.9.7.1. Configuring the cluster-wide proxy during installation	1813
10.9.7.2. Minimum resource requirements for cluster installation	1814
10.9.7.3. Tested instance types for IBM Cloud	1815
10.9.7.4. Sample customized install-config.yaml file for IBM Cloud	1816
10.9.8. Installing the OpenShift CLI	1819
Installing the OpenShift CLI on Linux	1819
Installing the OpenShift CLI on Windows	1819
Installing the OpenShift CLI on macOS	1820
10.9.9. Manually creating IAM	1820
10.9.10. Deploying the cluster	1823
10.9.11. Logging in to the cluster by using the CLI	1824
10.9.12. Post installation	1825
10.9.12.1. Disabling the default OperatorHub catalog sources	1825
10.9.12.2. Installing the policy resources into the cluster	1826
10.9.13. Telemetry access for OpenShift Container Platform	1826
10.9.14. Next steps	1827
10.10. INSTALLATION CONFIGURATION PARAMETERS FOR IBM CLOUD	1827
10.10.1. Available installation configuration parameters for IBM Cloud	1827
10.10.1.1. Required configuration parameters	1827
10.10.1.2. Network configuration parameters	1829
10.10.1.3. Optional configuration parameters	1831
10.10.1.4. Additional IBM Cloud configuration parameters	1837
10.11. UNINSTALLING A CLUSTER ON IBM CLOUD	1843
10.11.1. Removing a cluster that uses installer-provisioned infrastructure	1843

CHAPTER 11. INSTALLING ON NUTANIX	1845
11.1. PREPARING TO INSTALL ON NUTANIX	1845
11.1.1. Nutanix version requirements	1845
11.1.2. Environment requirements	1845
11.1.2.1. Required account privileges	1845
11.1.2.2. Cluster limits	1847
11.1.2.3. Cluster resources	1847
11.1.2.4. Networking requirements	1847
11.1.2.4.1. Required IP Addresses	1848
11.1.2.4.2. DNS records	1848
11.1.3. Configuring the Cloud Credential Operator utility	1848
11.2. FAULT TOLERANT DEPLOYMENTS USING MULTIPLE PRISM ELEMENTS	1850
11.2.1. Failure domain requirements	1850
11.2.2. Installation method and failure domain configuration	1851
11.3. INSTALLING A CLUSTER ON NUTANIX	1851
11.3.1. Prerequisites	1851
11.3.2. Internet access for OpenShift Container Platform	1852
11.3.3. Internet access for Prism Central	1852
11.3.4. Generating a key pair for cluster node SSH access	1852
11.3.5. Obtaining the installation program	1854
11.3.6. Adding Nutanix root CA certificates to your system trust	1855
11.3.7. Creating the installation configuration file	1855
11.3.7.1. Sample customized install-config.yaml file for Nutanix	1857
11.3.7.2. Configuring failure domains	1860
11.3.7.3. Configuring the cluster-wide proxy during installation	1862
11.3.8. Installing the OpenShift CLI	1863
Installing the OpenShift CLI on Linux	1864
Installing the OpenShift CLI on Windows	1864
Installing the OpenShift CLI on macOS	1865
11.3.9. Configuring IAM for Nutanix	1865
11.3.10. Adding config map and secret resources required for Nutanix CCM	1868
11.3.11. Services for a user-managed load balancer	1869
11.3.11.1. Configuring a user-managed load balancer	1872
11.3.12. Deploying the cluster	1879
11.3.13. Configuring the default storage container	1880
11.3.14. Telemetry access for OpenShift Container Platform	1880
11.3.15. Additional resources	1881
11.3.16. Next steps	1881
11.4. INSTALLING A CLUSTER ON NUTANIX IN A RESTRICTED NETWORK	1881
11.4.1. Prerequisites	1881
11.4.2. About installations in restricted networks	1882
11.4.2.1. Additional limits	1882
11.4.3. Generating a key pair for cluster node SSH access	1882
11.4.4. Adding Nutanix root CA certificates to your system trust	1884
11.4.5. Downloading the RHCOS cluster image	1884
11.4.6. Creating the installation configuration file	1885
11.4.6.1. Sample customized install-config.yaml file for Nutanix	1888
11.4.6.2. Configuring failure domains	1891
11.4.6.3. Configuring the cluster-wide proxy during installation	1893
11.4.7. Installing the OpenShift CLI	1895
Installing the OpenShift CLI on Linux	1895
Installing the OpenShift CLI on Windows	1895
Installing the OpenShift CLI on macOS	1896

11.4.8. Configuring IAM for Nutanix	1896
11.4.9. Deploying the cluster	1899
11.4.10. Post installation	1900
11.4.10.1. Disabling the default OperatorHub catalog sources	1901
11.4.10.2. Installing the policy resources into the cluster	1901
11.4.10.3. Configuring the default storage container	1902
11.4.11. Telemetry access for OpenShift Container Platform	1902
11.4.12. Additional resources	1902
11.4.13. Next steps	1902
11.5. INSTALLING A THREE-NODE CLUSTER ON NUTANIX	1902
11.5.1. Configuring a three-node cluster	1902
11.5.2. Next steps	1903
11.6. UNINSTALLING A CLUSTER ON NUTANIX	1903
11.6.1. Removing a cluster that uses installer-provisioned infrastructure	1903
11.7. INSTALLATION CONFIGURATION PARAMETERS FOR NUTANIX	1904
11.7.1. Available installation configuration parameters for Nutanix	1904
11.7.1.1. Required configuration parameters	1904
11.7.1.2. Network configuration parameters	1906
11.7.1.3. Optional configuration parameters	1907
11.7.1.4. Additional Nutanix configuration parameters	1913
CHAPTER 12. INSTALLING ON-PREMISE WITH ASSISTED INSTALLER	1919
12.1. INSTALLING AN ON-PREMISE CLUSTER USING THE ASSISTED INSTALLER	1919
12.1.1. Using the Assisted Installer	1919
12.1.2. API support for the Assisted Installer	1920
CHAPTER 13. INSTALLING AN ON-PREMISE CLUSTER WITH THE AGENT-BASED INSTALLER	1921
13.1. PREPARING TO INSTALL WITH THE AGENT-BASED INSTALLER	1921
13.1.1. About the Agent-based Installer	1921
13.1.2. Understanding Agent-based Installer	1921
13.1.2.1. Agent-based Installer workflow	1922
13.1.2.2. Recommended resources for topologies	1923
13.1.3. About FIPS compliance	1924
13.1.4. Configuring FIPS through the Agent-based Installer	1925
13.1.5. Host configuration	1925
13.1.5.1. Host roles	1926
13.1.5.2. About root device hints	1927
13.1.6. About networking	1928
13.1.6.1. DHCP	1928
13.1.6.2. Static networking	1928
13.1.7. Requirements for a cluster using the platform "none" option	1930
13.1.7.1. Platform "none" DNS requirements	1930
13.1.7.1.1. Example DNS configuration for platform "none" clusters	1932
13.1.7.2. Platform "none" Load balancing requirements	1934
13.1.7.2.1. Example load balancer configuration for platform "none" clusters	1936
13.1.8. Example: Bonds and VLAN interface node network configuration	1938
13.1.9. Example: Bonds and SR-IOV dual-nic node network configuration	1939
13.1.10. Sample install-config.yaml file for bare metal	1942
13.1.11. Validation checks before agent ISO creation	1945
13.1.11.1. ZTP manifests	1945
13.1.12. Next steps	1945
13.2. UNDERSTANDING DISCONNECTED INSTALLATION MIRRORING	1945
13.2.1. Mirroring images for a disconnected installation through the Agent-based Installer	1946

13.2.2. About mirroring the OpenShift Container Platform image repository for a disconnected registry	1946
13.2.2.1. Configuring the Agent-based Installer to use mirrored images	1947
13.3. INSTALLING AN OPENSIFT CONTAINER PLATFORM CLUSTER WITH THE AGENT-BASED INSTALLER	1948
13.3.1. Prerequisites	1948
13.3.2. Installing OpenShift Container Platform with the Agent-based Installer	1948
13.3.2.1. Downloading the Agent-based Installer	1948
13.3.2.2. Verifying the supported architecture for installing an Agent-based Installer cluster	1948
13.3.2.3. Creating the preferred configuration inputs	1949
13.3.2.4. Optional: Creating additional manifest files	1953
13.3.2.4.1. Creating a directory to contain additional manifests	1953
13.3.2.4.2. Disk partitioning	1953
13.3.2.5. Optional: Using ZTP manifests	1955
13.3.2.6. Optional: Encrypting the disk	1956
13.3.2.7. Creating and booting the agent image	1957
13.3.2.8. Adding IBM Z(R) agents with RHEL KVM	1958
13.3.2.9. Verifying that the current installation host can pull release images	1958
13.3.2.10. Tracking and verifying installation progress	1960
13.3.3. Sample GitOps ZTP custom resources	1962
13.3.4. Gathering log data from a failed Agent-based installation	1964
13.4. PREPARING PXE ASSETS FOR OPENSIFT CONTAINER PLATFORM	1965
13.4.1. Prerequisites	1966
13.4.2. Downloading the Agent-based Installer	1966
13.4.3. Creating the preferred configuration inputs	1966
13.4.4. Creating the PXE assets	1970
13.4.5. Manually adding IBM Z agents	1971
13.4.5.1. Adding IBM Z agents with z/VM	1971
13.4.5.2. Adding IBM Z(R) agents with RHEL KVM	1972
13.4.6. Additional resources	1973
13.5. PREPARING AN AGENT-BASED INSTALLED CLUSTER FOR THE MULTICLUSTER ENGINE FOR KUBERNETES OPERATOR	1973
13.5.1. Prerequisites	1973
13.5.2. Preparing an Agent-based cluster deployment for the multicluster engine for Kubernetes Operator while disconnected	1973
13.5.3. Preparing an Agent-based cluster deployment for the multicluster engine for Kubernetes Operator while connected	1976
13.6. INSTALLATION CONFIGURATION PARAMETERS FOR THE AGENT-BASED INSTALLER	1980
13.6.1. Available installation configuration parameters	1980
13.6.1.1. Required configuration parameters	1981
13.6.1.2. Network configuration parameters	1982
13.6.1.3. Optional configuration parameters	1984
13.6.1.4. Additional bare metal configuration parameters for the Agent-based Installer	1990
13.6.1.5. Additional VMware vSphere configuration parameters	1993
13.6.1.6. Deprecated VMware vSphere configuration parameters	1997
13.6.2. Available Agent configuration parameters	1998
13.6.2.1. Required configuration parameters	1998
13.6.2.2. Optional configuration parameters	1999
CHAPTER 14. INSTALLING ON A SINGLE NODE	2002
14.1. PREPARING TO INSTALL ON A SINGLE NODE	2002
14.1.1. Prerequisites	2002
14.1.2. About OpenShift on a single node	2002
14.1.3. Requirements for installing OpenShift on a single node	2002

14.2. INSTALLING OPENSIFT ON A SINGLE NODE	2004
14.2.1. Installing single-node OpenShift using the Assisted Installer	2004
14.2.1.1. Generating the discovery ISO with the Assisted Installer	2004
14.2.1.2. Installing single-node OpenShift with the Assisted Installer	2005
14.2.2. Installing single-node OpenShift manually	2005
14.2.2.1. Generating the installation ISO with coreos-installer	2006
14.2.2.2. Monitoring the cluster installation using openshift-install	2008
14.2.3. Installing single-node OpenShift on cloud providers	2009
14.2.3.1. Additional requirements for installing single-node OpenShift on a cloud provider	2009
14.2.3.2. Supported cloud providers for single-node OpenShift	2009
14.2.3.3. Installing single-node OpenShift on AWS	2010
14.2.3.4. Installing single-node OpenShift on Azure	2010
14.2.3.5. Installing single-node OpenShift on GCP	2010
14.2.4. Creating a bootable ISO image on a USB drive	2010
14.2.5. Booting from an HTTP-hosted ISO image using the Redfish API	2010
14.2.6. Creating a custom live RHCOS ISO for remote server access	2012
14.2.7. Installing single-node OpenShift with IBM Z and IBM LinuxONE	2013
Hardware requirements	2013
14.2.7.1. Installing single-node OpenShift with z/VM on IBM Z and IBM LinuxONE	2014
14.2.7.2. Installing single-node OpenShift with RHEL KVM on IBM Z and IBM LinuxONE	2019
14.2.7.3. Installing single-node OpenShift in an LPAR on IBM Z and IBM LinuxONE	2022
14.2.8. Installing single-node OpenShift with IBM Power	2026
Hardware requirements	2026
14.2.8.1. Setting up basion for single-node OpenShift with IBM Power	2027
14.2.8.2. Installing single-node OpenShift with IBM Power	2030
CHAPTER 15. INSTALLING ON BARE METAL	2031
15.1. PREPARING FOR BARE METAL CLUSTER INSTALLATION	2031
15.1.1. Prerequisites	2031
15.1.2. Planning a bare metal cluster for OpenShift Virtualization	2031
15.1.3. NIC partitioning for SR-IOV devices (Technology Preview)	2031
15.1.4. Choosing a method to install OpenShift Container Platform on bare metal	2032
15.1.4.1. Installing a cluster on installer-provisioned infrastructure	2033
15.1.4.2. Installing a cluster on user-provisioned infrastructure	2033
15.2. INSTALLING A USER-PROVISIONED CLUSTER ON BARE METAL	2033
15.2.1. Prerequisites	2034
15.2.2. Internet access for OpenShift Container Platform	2034
15.2.3. Requirements for a cluster with user-provisioned infrastructure	2034
15.2.3.1. Required machines for cluster installation	2035
15.2.3.2. Minimum resource requirements for cluster installation	2035
15.2.3.3. Certificate signing requests management	2036
15.2.3.4. Requirements for baremetal clusters on vSphere	2037
15.2.3.5. Networking requirements for user-provisioned infrastructure	2037
15.2.3.5.1. Setting the cluster node hostnames through DHCP	2038
15.2.3.5.2. Network connectivity requirements	2038
NTP configuration for user-provisioned infrastructure	2039
15.2.3.6. User-provisioned DNS requirements	2039
15.2.3.6.1. Example DNS configuration for user-provisioned clusters	2041
15.2.3.7. Load balancing requirements for user-provisioned infrastructure	2043
15.2.3.7.1. Example load balancer configuration for user-provisioned clusters	2045
15.2.4. Preparing the user-provisioned infrastructure	2047
15.2.5. Validating DNS resolution for user-provisioned infrastructure	2049
15.2.6. Generating a key pair for cluster node SSH access	2052

15.2.7. Obtaining the installation program	2053
15.2.8. Installing the OpenShift CLI	2054
Installing the OpenShift CLI on Linux	2054
Installing the OpenShift CLI on Windows	2055
Installing the OpenShift CLI on macOS	2055
15.2.9. Manually creating the installation configuration file	2056
15.2.9.1. Sample install-config.yaml file for bare metal	2057
15.2.9.2. Configuring the cluster-wide proxy during installation	2060
15.2.9.3. Configuring a three-node cluster	2061
15.2.10. Creating the Kubernetes manifest and Ignition config files	2062
15.2.11. Installing RHCOS and starting the OpenShift Container Platform bootstrap process	2064
15.2.11.1. Installing RHCOS by using an ISO image	2065
15.2.11.2. Installing RHCOS by using PXE or iPXE booting	2068
15.2.11.3. Advanced RHCOS installation configuration	2073
15.2.11.3.1. Using advanced networking options for PXE and ISO installations	2073
15.2.11.3.2. Disk partitioning	2074
15.2.11.3.2.1. Creating a separate /var partition	2075
15.2.11.3.2.2. Retaining existing partitions	2077
15.2.11.3.3. Identifying Ignition configs	2078
15.2.11.3.4. Default console configuration	2078
15.2.11.3.5. Enabling the serial console for PXE and ISO installations	2079
15.2.11.3.6. Customizing a live RHCOS ISO or PXE install	2080
15.2.11.3.7. Customizing a live RHCOS ISO image	2080
15.2.11.3.7.1. Modifying a live install ISO image to enable the serial console	2081
15.2.11.3.7.2. Modifying a live install ISO image to use a custom certificate authority	2082
15.2.11.3.7.3. Modifying a live install ISO image with customized network settings	2082
15.2.11.3.7.4. Customizing a live install ISO image for an iSCSI boot device	2084
15.2.11.3.7.5. Customizing a live install ISO image for an iSCSI boot device with iBFT	2085
15.2.11.3.8. Customizing a live RHCOS PXE environment	2085
15.2.11.3.8.1. Modifying a live install PXE environment to enable the serial console	2086
15.2.11.3.8.2. Modifying a live install PXE environment to use a custom certificate authority	2087
15.2.11.3.8.3. Modifying a live install PXE environment with customized network settings	2087
15.2.11.3.8.4. Customizing a live install PXE environment for an iSCSI boot device	2089
15.2.11.3.8.5. Customizing a live install PXE environment for an iSCSI boot device with iBFT	2090
15.2.11.3.9. Advanced RHCOS installation reference	2091
15.2.11.3.9.1. Networking and bonding options for ISO installations	2091
Configuring DHCP or static IP addresses	2091
Configuring an IP address without a static hostname	2092
Specifying multiple network interfaces	2092
Configuring default gateway and route	2092
Disabling DHCP on a single interface	2093
Combining DHCP and static IP configurations	2093
Configuring VLANs on individual interfaces	2093
Providing multiple DNS servers	2093
Bonding multiple network interfaces to a single interface	2093
Bonding multiple SR-IOV network interfaces to a dual port NIC interface	2094
Using network teaming	2095
15.2.11.3.9.2. coreos-installer options for ISO and PXE installations	2095
15.2.11.3.9.3. coreos.inst boot options for ISO or PXE installations	2099
15.2.11.4. Enabling multipathing with kernel arguments on RHCOS	2101
15.2.11.5. Installing RHCOS manually on an iSCSI boot device	2102
15.2.11.6. Installing RHCOS on an iSCSI boot device using iBFT	2103
15.2.12. Waiting for the bootstrap process to complete	2105

15.2.13. Logging in to the cluster by using the CLI	2106
15.2.14. Approving the certificate signing requests for your machines	2106
15.2.15. Initial Operator configuration	2109
15.2.15.1. Image registry removed during installation	2110
15.2.15.2. Image registry storage configuration	2110
15.2.15.2.1. Configuring registry storage for bare metal and other manual installations	2111
15.2.15.2.2. Configuring storage for the image registry in non-production clusters	2112
15.2.15.2.3. Configuring block registry storage for bare metal	2113
15.2.16. Completing installation on user-provisioned infrastructure	2114
15.2.17. Telemetry access for OpenShift Container Platform	2117
15.2.18. Next steps	2117
15.3. INSTALLING A USER-PROVISIONED BARE METAL CLUSTER WITH NETWORK CUSTOMIZATIONS	2117
15.3.1. Prerequisites	2117
15.3.2. Internet access for OpenShift Container Platform	2117
15.3.3. Requirements for a cluster with user-provisioned infrastructure	2118
15.3.3.1. Required machines for cluster installation	2118
15.3.3.2. Minimum resource requirements for cluster installation	2119
15.3.3.3. Certificate signing requests management	2120
15.3.3.4. Networking requirements for user-provisioned infrastructure	2120
15.3.3.4.1. Setting the cluster node hostnames through DHCP	2121
15.3.3.4.2. Network connectivity requirements	2121
NTP configuration for user-provisioned infrastructure	2122
15.3.3.5. User-provisioned DNS requirements	2123
15.3.3.5.1. Example DNS configuration for user-provisioned clusters	2125
15.3.3.6. Load balancing requirements for user-provisioned infrastructure	2127
15.3.3.6.1. Example load balancer configuration for user-provisioned clusters	2129
15.3.4. Preparing the user-provisioned infrastructure	2131
15.3.5. Validating DNS resolution for user-provisioned infrastructure	2133
15.3.6. Generating a key pair for cluster node SSH access	2135
15.3.7. Obtaining the installation program	2137
15.3.8. Installing the OpenShift CLI	2138
Installing the OpenShift CLI on Linux	2138
Installing the OpenShift CLI on Windows	2139
Installing the OpenShift CLI on macOS	2139
15.3.9. Manually creating the installation configuration file	2140
15.3.9.1. Sample install-config.yaml file for bare metal	2141
15.3.10. Network configuration phases	2143
15.3.11. Specifying advanced network configuration	2144
15.3.12. Cluster Network Operator configuration	2145
15.3.12.1. Cluster Network Operator configuration object	2145
defaultNetwork object configuration	2146
Configuration for the OVN-Kubernetes network plugin	2147
kubeProxyConfig object configuration (OpenShiftSDN container network interface only)	2151
15.3.13. Creating the Ignition config files	2152
15.3.14. Installing RHCOS and starting the OpenShift Container Platform bootstrap process	2153
15.3.14.1. Installing RHCOS by using an ISO image	2154
15.3.14.2. Installing RHCOS by using PXE or iPXE booting	2157
15.3.14.3. Advanced RHCOS installation configuration	2162
15.3.14.3.1. Using advanced networking options for PXE and ISO installations	2162
15.3.14.3.2. Disk partitioning	2163
15.3.14.3.2.1. Creating a separate /var partition	2164
15.3.14.3.2.2. Retaining existing partitions	2166
15.3.14.3.3. Identifying Ignition configs	2167

15.3.14.3.4. Default console configuration	2167
15.3.14.3.5. Enabling the serial console for PXE and ISO installations	2168
15.3.14.3.6. Customizing a live RHCOS ISO or PXE install	2169
15.3.14.3.7. Customizing a live RHCOS ISO image	2169
15.3.14.3.7.1. Modifying a live install ISO image to enable the serial console	2170
15.3.14.3.7.2. Modifying a live install ISO image to use a custom certificate authority	2170
15.3.14.3.7.3. Modifying a live install ISO image with customized network settings	2171
15.3.14.3.7.4. Customizing a live install ISO image for an iSCSI boot device	2173
15.3.14.3.7.5. Customizing a live install ISO image for an iSCSI boot device with iBFT	2173
15.3.14.3.8. Customizing a live RHCOS PXE environment	2174
15.3.14.3.8.1. Modifying a live install PXE environment to enable the serial console	2175
15.3.14.3.8.2. Modifying a live install PXE environment to use a custom certificate authority	2176
15.3.14.3.8.3. Modifying a live install PXE environment with customized network settings	2176
15.3.14.3.8.4. Customizing a live install PXE environment for an iSCSI boot device	2178
15.3.14.3.8.5. Customizing a live install PXE environment for an iSCSI boot device with iBFT	2179
15.3.14.3.9. Advanced RHCOS installation reference	2180
15.3.14.3.9.1. Networking and bonding options for ISO installations	2180
Configuring DHCP or static IP addresses	2180
Configuring an IP address without a static hostname	2181
Specifying multiple network interfaces	2181
Configuring default gateway and route	2181
Disabling DHCP on a single interface	2181
Combining DHCP and static IP configurations	2182
Configuring VLANs on individual interfaces	2182
Providing multiple DNS servers	2182
Bonding multiple network interfaces to a single interface	2182
Bonding multiple SR-IOV network interfaces to a dual port NIC interface	2183
Using network teaming	2184
15.3.14.3.9.2. coreos-installer options for ISO and PXE installations	2184
15.3.14.3.9.3. coreos.inst boot options for ISO or PXE installations	2188
15.3.14.4. Enabling multipathing with kernel arguments on RHCOS	2190
15.3.14.5. Installing RHCOS manually on an iSCSI boot device	2191
15.3.14.6. Installing RHCOS on an iSCSI boot device using iBFT	2192
15.3.15. Waiting for the bootstrap process to complete	2193
15.3.16. Logging in to the cluster by using the CLI	2194
15.3.17. Approving the certificate signing requests for your machines	2195
15.3.18. Initial Operator configuration	2197
15.3.18.1. Image registry removed during installation	2199
15.3.18.2. Image registry storage configuration	2199
15.3.18.3. Configuring block registry storage for bare metal	2199
15.3.19. Completing installation on user-provisioned infrastructure	2200
15.3.20. Telemetry access for OpenShift Container Platform	2203
15.3.21. Next steps	2203
15.4. INSTALLING A USER-PROVISIONED BARE METAL CLUSTER ON A RESTRICTED NETWORK	2203
15.4.1. Prerequisites	2203
15.4.2. About installations in restricted networks	2204
15.4.2.1. Additional limits	2204
15.4.3. Internet access for OpenShift Container Platform	2204
15.4.4. Requirements for a cluster with user-provisioned infrastructure	2205
15.4.4.1. Required machines for cluster installation	2205
15.4.4.2. Minimum resource requirements for cluster installation	2206
15.4.4.3. Certificate signing requests management	2207
15.4.4.4. Networking requirements for user-provisioned infrastructure	2207

15.4.4.4.1. Setting the cluster node hostnames through DHCP	2208
15.4.4.4.2. Network connectivity requirements	2208
NTP configuration for user-provisioned infrastructure	2209
15.4.4.5. User-provisioned DNS requirements	2209
15.4.4.5.1. Example DNS configuration for user-provisioned clusters	2211
15.4.4.6. Load balancing requirements for user-provisioned infrastructure	2214
15.4.4.6.1. Example load balancer configuration for user-provisioned clusters	2216
15.4.5. Preparing the user-provisioned infrastructure	2217
15.4.6. Validating DNS resolution for user-provisioned infrastructure	2220
15.4.7. Generating a key pair for cluster node SSH access	2222
15.4.8. Manually creating the installation configuration file	2224
15.4.8.1. Sample install-config.yaml file for bare metal	2225
15.4.8.2. Configuring the cluster-wide proxy during installation	2228
15.4.8.3. Configuring a three-node cluster	2230
15.4.9. Creating the Kubernetes manifest and Ignition config files	2231
15.4.10. Configuring chrony time service	2233
15.4.11. Installing RHCOS and starting the OpenShift Container Platform bootstrap process	2234
15.4.11.1. Installing RHCOS by using an ISO image	2235
15.4.11.2. Installing RHCOS by using PXE or iPXE booting	2238
15.4.11.3. Advanced RHCOS installation configuration	2243
15.4.11.3.1. Using advanced networking options for PXE and ISO installations	2243
15.4.11.3.2. Disk partitioning	2244
15.4.11.3.2.1. Creating a separate /var partition	2245
15.4.11.3.2.2. Retaining existing partitions	2247
15.4.11.3.3. Identifying Ignition configs	2248
15.4.11.3.4. Default console configuration	2248
15.4.11.3.5. Enabling the serial console for PXE and ISO installations	2249
15.4.11.3.6. Customizing a live RHCOS ISO or PXE install	2250
15.4.11.3.7. Customizing a live RHCOS ISO image	2250
15.4.11.3.7.1. Modifying a live install ISO image to enable the serial console	2251
15.4.11.3.7.2. Modifying a live install ISO image to use a custom certificate authority	2251
15.4.11.3.7.3. Modifying a live install ISO image with customized network settings	2252
15.4.11.3.7.4. Customizing a live install ISO image for an iSCSI boot device	2254
15.4.11.3.7.5. Customizing a live install ISO image for an iSCSI boot device with iBFT	2254
15.4.11.3.8. Customizing a live RHCOS PXE environment	2255
15.4.11.3.8.1. Modifying a live install PXE environment to enable the serial console	2256
15.4.11.3.8.2. Modifying a live install PXE environment to use a custom certificate authority	2257
15.4.11.3.8.3. Modifying a live install PXE environment with customized network settings	2257
15.4.11.3.8.4. Customizing a live install PXE environment for an iSCSI boot device	2259
15.4.11.3.8.5. Customizing a live install PXE environment for an iSCSI boot device with iBFT	2260
15.4.11.3.9. Advanced RHCOS installation reference	2261
15.4.11.3.9.1. Networking and bonding options for ISO installations	2261
Configuring DHCP or static IP addresses	2261
Configuring an IP address without a static hostname	2262
Specifying multiple network interfaces	2262
Configuring default gateway and route	2262
Disabling DHCP on a single interface	2262
Combining DHCP and static IP configurations	2263
Configuring VLANs on individual interfaces	2263
Providing multiple DNS servers	2263
Bonding multiple network interfaces to a single interface	2263
Bonding multiple SR-IOV network interfaces to a dual port NIC interface	2264
Using network teaming	2265

15.4.11.3.9.2. coreos-installer options for ISO and PXE installations	2265
15.4.11.3.9.3. coreos.inst boot options for ISO or PXE installations	2269
15.4.11.4. Enabling multipathing with kernel arguments on RHCOS	2271
15.4.11.5. Installing RHCOS manually on an iSCSI boot device	2272
15.4.11.6. Installing RHCOS on an iSCSI boot device using iBFT	2273
15.4.12. Waiting for the bootstrap process to complete	2274
15.4.13. Logging in to the cluster by using the CLI	2275
15.4.14. Approving the certificate signing requests for your machines	2276
15.4.15. Initial Operator configuration	2278
15.4.15.1. Disabling the default OperatorHub catalog sources	2280
15.4.15.2. Image registry storage configuration	2280
15.4.15.2.1. Changing the image registry's management state	2280
15.4.15.2.2. Configuring registry storage for bare metal and other manual installations	2280
15.4.15.2.3. Configuring storage for the image registry in non-production clusters	2282
15.4.15.2.4. Configuring block registry storage for bare metal	2283
15.4.16. Completing installation on user-provisioned infrastructure	2284
15.4.17. Telemetry access for OpenShift Container Platform	2287
15.4.18. Next steps	2287
15.5. SCALING A USER-PROVISIONED CLUSTER WITH THE BARE METAL OPERATOR	2287
15.5.1. About scaling a user-provisioned cluster with the Bare Metal Operator	2287
15.5.1.1. Prerequisites for scaling a user-provisioned cluster	2287
15.5.1.2. Limitations for scaling a user-provisioned cluster	2288
15.5.2. Configuring a provisioning resource to scale user-provisioned clusters	2288
15.5.3. Provisioning new hosts in a user-provisioned cluster by using the BMO	2289
15.5.4. Optional: Managing existing hosts in a user-provisioned cluster by using the BMO	2291
15.5.5. Removing hosts from a user-provisioned cluster by using the BMO	2293
15.6. INSTALLATION CONFIGURATION PARAMETERS FOR BARE METAL	2294
15.6.1. Available installation configuration parameters for bare metal	2294
15.6.1.1. Required configuration parameters	2294
15.6.1.2. Network configuration parameters	2296
15.6.1.3. Optional configuration parameters	2298
CHAPTER 16. DEPLOYING INSTALLER-PROVISIONED CLUSTERS ON BARE METAL	2305
16.1. OVERVIEW	2305
16.2. PREREQUISITES	2306
16.2.1. Node requirements	2307
16.2.2. Planning a bare metal cluster for OpenShift Virtualization	2308
16.2.3. Firmware requirements for installing with virtual media	2309
16.2.4. Network requirements	2310
16.2.4.1. Ensuring required ports are open	2310
16.2.4.2. Increase the network MTU	2312
16.2.4.3. Configuring NICs	2312
16.2.4.4. DNS requirements	2313
16.2.4.5. Dynamic Host Configuration Protocol (DHCP) requirements	2314
16.2.4.6. Reserving IP addresses for nodes with the DHCP server	2314
16.2.4.7. Provisioner node requirements	2315
16.2.4.8. Network Time Protocol (NTP)	2316
16.2.4.9. Port access for the out-of-band management IP address	2316
16.2.5. Configuring nodes	2316
Configuring nodes when using the provisioning network	2316
Configuring nodes without the provisioning network	2317
Configuring nodes for Secure Boot manually	2317
16.2.6. Out-of-band management	2318

16.2.7. Required data for installation	2318
16.2.8. Validation checklist for nodes	2319
16.3. SETTING UP THE ENVIRONMENT FOR AN OPENSIFT INSTALLATION	2319
16.3.1. Installing RHEL on the provisioner node	2320
16.3.2. Preparing the provisioner node for OpenShift Container Platform installation	2320
16.3.3. Checking NTP server synchronization	2321
16.3.4. Configuring networking	2322
16.3.5. Establishing communication between subnets	2324
16.3.6. Retrieving the OpenShift Container Platform installer	2327
16.3.7. Extracting the OpenShift Container Platform installer	2327
16.3.8. Optional: Creating an RHCOS images cache	2328
16.3.9. Services for a user-managed load balancer	2330
16.3.9.1. Configuring a user-managed load balancer	2333
16.3.10. Setting the cluster node hostnames through DHCP	2340
16.3.11. Configuring the install-config.yaml file	2340
16.3.11.1. Configuring the install-config.yaml file	2340
16.3.11.2. Additional install-config parameters	2344
Hosts	2348
16.3.11.3. BMC addressing	2349
IPMI	2349
Redfish network boot	2350
16.3.11.4. Verifying support for Redfish APIs	2350
16.3.11.5. BMC addressing for Dell iDRAC	2352
BMC address formats for Dell iDRAC	2352
Redfish virtual media for Dell iDRAC	2353
Redfish network boot for iDRAC	2354
16.3.11.6. BMC addressing for HPE iLO	2354
Redfish virtual media for HPE iLO	2355
Redfish network boot for HPE iLO	2356
16.3.11.7. BMC addressing for Fujitsu iRMC	2356
16.3.11.8. Root device hints	2357
16.3.11.9. Optional: Setting proxy settings	2359
16.3.11.10. Optional: Deploying with no provisioning network	2359
16.3.11.11. Optional: Deploying with dual-stack networking	2360
16.3.11.12. Optional: Configuring host network interfaces	2361
16.3.11.13. Configuring host network interfaces for subnets	2363
16.3.11.14. Optional: Configuring address generation modes for SLAAC in dual-stack networks	2364
16.3.11.15. Optional: Configuring host network interfaces for dual port NIC	2365
16.3.11.16. Configuring multiple cluster nodes	2368
16.3.11.17. Optional: Configuring managed Secure Boot	2369
16.3.12. Manifest configuration files	2370
16.3.12.1. Creating the OpenShift Container Platform manifests	2370
16.3.12.2. Optional: Configuring NTP for disconnected clusters	2370
16.3.12.3. Configuring network components to run on the control plane	2373
16.3.12.4. Optional: Deploying routers on compute nodes	2375
16.3.12.5. Optional: Configuring the BIOS	2376
16.3.12.6. Optional: Configuring the RAID	2377
16.3.12.7. Optional: Configuring storage on nodes	2378
16.3.13. Creating a disconnected registry	2379
Prerequisites	2379
16.3.13.1. Preparing the registry node to host the mirrored registry	2379
16.3.13.2. Mirroring the OpenShift Container Platform image repository for a disconnected registry	2380
16.3.13.3. Modify the install-config.yaml file to use the disconnected registry	2383

16.3.14. Validation checklist for installation	2384
16.3.15. Deploying the cluster via the OpenShift Container Platform installer	2384
16.3.16. Following the installation	2384
16.3.17. Verifying static IP address configuration	2384
16.3.18. Preparing to reinstall a cluster on bare metal	2385
16.3.19. Additional resources	2385
16.4. INSTALLER-PROVISIONED POSTINSTALLATION CONFIGURATION	2385
16.4.1. Optional: Configuring NTP for disconnected clusters	2385
16.4.2. Enabling a provisioning network after installation	2388
16.5. EXPANDING THE CLUSTER	2390
16.5.1. Preparing the bare metal node	2390
16.5.2. Replacing a bare-metal control plane node	2395
16.5.3. Preparing to deploy with Virtual Media on the baremetal network	2399
16.5.4. Diagnosing a duplicate MAC address when provisioning a new host in the cluster	2401
16.5.5. Provisioning the bare metal node	2402
16.6. TROUBLESHOOTING	2403
16.6.1. Troubleshooting the installation program workflow	2403
16.6.2. Troubleshooting install-config.yaml	2406
16.6.3. Troubleshooting bootstrap VM issues	2406
16.6.3.1. Bootstrap VM cannot boot up the cluster nodes	2408
16.6.3.2. Inspecting logs	2408
16.6.4. Investigating an unavailable Kubernetes API	2409
16.6.5. Troubleshooting a failure to initialize the cluster	2410
16.6.6. Troubleshooting a failure to fetch the console URL	2414
16.6.7. Troubleshooting a failure to add the ingress certificate to kubeconfig	2415
16.6.8. Troubleshooting SSH access to cluster nodes	2415
16.6.9. Cluster nodes will not PXE boot	2416
16.6.10. Installing creates no worker nodes	2416
16.6.11. Troubleshooting the Cluster Network Operator	2417
16.6.12. Unable to discover new bare metal hosts using the BMC	2417
16.6.13. Troubleshooting worker nodes that cannot join the cluster	2418
16.6.14. Cleaning up previous installations	2419
16.6.15. Issues with creating the registry	2419
16.6.16. Miscellaneous issues	2420
16.6.16.1. Addressing the runtime network not ready error	2420
16.6.16.2. Addressing the "No disk found with matching rootDeviceHints" error message	2421
16.6.16.3. Cluster nodes not getting the correct IPv6 address over DHCP	2421
16.6.16.4. Cluster nodes not getting the correct hostname over DHCP	2421
16.6.16.5. Routes do not reach endpoints	2423
16.6.16.6. Failed Ignition during Firstboot	2424
16.6.16.7. NTP out of sync	2424
16.6.17. Reviewing the installation	2426
CHAPTER 17. INSTALLING IBM CLOUD BARE METAL (CLASSIC)	2428
17.1. PREREQUISITES	2428
17.1.1. Setting up IBM Cloud Bare Metal (Classic) infrastructure	2428
Use one data center per cluster	2428
Create public and private VLANs	2428
Ensure subnets have sufficient IP addresses	2428
Configuring NICs	2429
Configuring canonical names	2430
Creating DNS entries	2430
Network Time Protocol (NTP)	2431

Configure a DHCP server	2431
Ensure BMC access privileges	2431
Create bare metal servers	2432
17.2. SETTING UP THE ENVIRONMENT FOR AN OPENSIFT CONTAINER PLATFORM INSTALLATION	2432
17.2.1. Preparing the provisioner node on IBM Cloud(R) Bare Metal (Classic) infrastructure	2432
17.2.2. Configuring the public subnet	2436
17.2.3. Retrieving the OpenShift Container Platform installer	2438
17.2.4. Extracting the OpenShift Container Platform installer	2439
17.2.5. Configuring the install-config.yaml file	2439
17.2.6. Additional install-config parameters	2441
Hosts	2445
17.2.7. Root device hints	2446
17.2.8. Creating the OpenShift Container Platform manifests	2447
17.2.9. Deploying the cluster via the OpenShift Container Platform installer	2448
17.2.10. Following the installation	2448
CHAPTER 18. INSTALLING ON IBM Z AND IBM LINUXONE	2449
18.1. PREPARING TO INSTALL ON IBM Z AND IBM LINUXONE	2449
18.1.1. Prerequisites	2449
18.1.2. Choosing a method to install OpenShift Container Platform on IBM Z or IBM LinuxONE	2449
18.1.2.1. User-provisioned infrastructure installation of OpenShift Container Platform on IBM Z	2450
18.2. INSTALLING A CLUSTER WITH Z/VM ON IBM Z AND IBM LINUXONE	2451
18.2.1. Prerequisites	2451
18.2.2. Internet access for OpenShift Container Platform	2452
18.2.3. Requirements for a cluster with user-provisioned infrastructure	2452
18.2.3.1. Required machines for cluster installation	2452
18.2.3.2. Minimum resource requirements for cluster installation	2453
18.2.3.3. Minimum IBM Z system environment	2454
Hardware requirements	2454
Operating system requirements	2454
IBM Z network connectivity requirements	2454
Disk storage	2454
Storage / Main Memory	2455
18.2.3.4. Preferred IBM Z system environment	2455
Hardware requirements	2455
Operating system requirements	2455
IBM Z network connectivity requirements	2455
Disk storage	2455
Storage / Main Memory	2456
18.2.3.5. Certificate signing requests management	2456
18.2.3.6. Networking requirements for user-provisioned infrastructure	2456
18.2.3.6.1. Network connectivity requirements	2456
NTP configuration for user-provisioned infrastructure	2458
18.2.3.7. User-provisioned DNS requirements	2458
18.2.3.7.1. Example DNS configuration for user-provisioned clusters	2460
18.2.3.8. Load balancing requirements for user-provisioned infrastructure	2462
18.2.3.8.1. Example load balancer configuration for user-provisioned clusters	2464
18.2.4. Preparing the user-provisioned infrastructure	2465
18.2.5. Validating DNS resolution for user-provisioned infrastructure	2467
18.2.6. Generating a key pair for cluster node SSH access	2469
18.2.7. Obtaining the installation program	2471
18.2.8. Installing the OpenShift CLI	2472
Installing the OpenShift CLI on Linux	2472

Installing the OpenShift CLI on Windows	2472
Installing the OpenShift CLI on macOS	2473
18.2.9. Manually creating the installation configuration file	2473
18.2.9.1. Sample install-config.yaml file for IBM Z	2474
18.2.9.2. Configuring the cluster-wide proxy during installation	2477
18.2.9.3. Configuring a three-node cluster	2479
18.2.10. Cluster Network Operator configuration	2480
18.2.10.1. Cluster Network Operator configuration object	2480
defaultNetwork object configuration	2481
Configuration for the OVN-Kubernetes network plugin	2481
kubeProxyConfig object configuration (OpenShiftSDN container network interface only)	2486
18.2.11. Creating the Kubernetes manifest and Ignition config files	2486
18.2.12. Configuring NBDE with static IP in an IBM Z or IBM LinuxONE environment	2488
18.2.13. Installing RHCOS and starting the OpenShift Container Platform bootstrap process	2490
18.2.13.1. Advanced RHCOS installation reference	2493
18.2.13.1.1. Networking and bonding options for ISO installations	2494
Configuring DHCP or static IP addresses	2494
Configuring an IP address without a static hostname	2495
Specifying multiple network interfaces	2495
Configuring default gateway and route	2495
Disabling DHCP on a single interface	2495
Combining DHCP and static IP configurations	2496
Configuring VLANs on individual interfaces	2496
Providing multiple DNS servers	2496
Bonding multiple network interfaces to a single interface	2496
Bonding multiple network interfaces to a single interface	2497
Using network teaming	2497
18.2.14. Waiting for the bootstrap process to complete	2497
18.2.15. Logging in to the cluster by using the CLI	2498
18.2.16. Approving the certificate signing requests for your machines	2499
18.2.17. Initial Operator configuration	2501
18.2.17.1. Image registry storage configuration	2502
18.2.17.1.1. Configuring registry storage for IBM Z	2502
18.2.17.1.2. Configuring storage for the image registry in non-production clusters	2504
18.2.18. Completing installation on user-provisioned infrastructure	2505
18.2.19. Telemetry access for OpenShift Container Platform	2507
18.2.20. Next steps	2508
18.3. INSTALLING A CLUSTER WITH Z/VM ON IBM Z AND IBM LINUXONE IN A RESTRICTED NETWORK	2508
18.3.1. Prerequisites	2508
18.3.2. About installations in restricted networks	2509
18.3.2.1. Additional limits	2509
18.3.3. Internet access for OpenShift Container Platform	2509
18.3.4. Requirements for a cluster with user-provisioned infrastructure	2509
18.3.4.1. Required machines for cluster installation	2510
18.3.4.2. Minimum resource requirements for cluster installation	2510
18.3.4.3. Minimum IBM Z system environment	2511
Hardware requirements	2511
Operating system requirements	2512
IBM Z network connectivity requirements	2512
Disk storage	2512
Storage / Main Memory	2512
18.3.4.4. Preferred IBM Z system environment	2512

Hardware requirements	2512
Operating system requirements	2513
IBM Z network connectivity requirements	2513
Disk storage	2513
Storage / Main Memory	2513
18.3.4.5. Certificate signing requests management	2513
18.3.4.6. Networking requirements for user-provisioned infrastructure	2514
18.3.4.6.1. Setting the cluster node hostnames through DHCP	2514
18.3.4.6.2. Network connectivity requirements	2514
NTP configuration for user-provisioned infrastructure	2515
18.3.4.7. User-provisioned DNS requirements	2516
18.3.4.7.1. Example DNS configuration for user-provisioned clusters	2518
18.3.4.8. Load balancing requirements for user-provisioned infrastructure	2520
18.3.4.8.1. Example load balancer configuration for user-provisioned clusters	2522
18.3.5. Preparing the user-provisioned infrastructure	2523
18.3.6. Validating DNS resolution for user-provisioned infrastructure	2525
18.3.7. Generating a key pair for cluster node SSH access	2527
18.3.8. Manually creating the installation configuration file	2529
18.3.8.1. Sample install-config.yaml file for IBM Z	2530
18.3.8.2. Configuring the cluster-wide proxy during installation	2533
18.3.8.3. Configuring a three-node cluster	2534
18.3.9. Cluster Network Operator configuration	2536
18.3.9.1. Cluster Network Operator configuration object	2536
defaultNetwork object configuration	2537
Configuration for the OVN-Kubernetes network plugin	2537
kubeProxyConfig object configuration (OpenShiftSDN container network interface only)	2542
18.3.10. Creating the Kubernetes manifest and Ignition config files	2542
18.3.11. Configuring NBDE with static IP in an IBM Z or IBM LinuxONE environment	2544
18.3.12. Installing RHCOS and starting the OpenShift Container Platform bootstrap process	2546
18.3.12.1. Advanced RHCOS installation reference	2550
18.3.12.1.1. Networking and bonding options for ISO installations	2550
Configuring DHCP or static IP addresses	2550
Configuring an IP address without a static hostname	2551
Specifying multiple network interfaces	2551
Configuring default gateway and route	2551
Disabling DHCP on a single interface	2551
Combining DHCP and static IP configurations	2552
Configuring VLANs on individual interfaces	2552
Providing multiple DNS servers	2552
Bonding multiple network interfaces to a single interface	2552
Bonding multiple network interfaces to a single interface	2553
Using network teaming	2553
18.3.13. Waiting for the bootstrap process to complete	2553
18.3.14. Logging in to the cluster by using the CLI	2554
18.3.15. Approving the certificate signing requests for your machines	2555
18.3.16. Initial Operator configuration	2557
18.3.16.1. Disabling the default OperatorHub catalog sources	2558
18.3.16.2. Image registry storage configuration	2559
18.3.16.2.1. Configuring registry storage for IBM Z	2559
18.3.16.2.2. Configuring storage for the image registry in non-production clusters	2561
18.3.17. Completing installation on user-provisioned infrastructure	2561
18.3.18. Next steps	2564
18.4. INSTALLING A CLUSTER WITH RHEL KVM ON IBM Z AND IBM LINUXONE	2564

18.4.1. Prerequisites	2565
18.4.2. Internet access for OpenShift Container Platform	2565
18.4.3. Machine requirements for a cluster with user-provisioned infrastructure	2566
18.4.3.1. Required machines	2566
18.4.3.2. Network connectivity requirements	2566
18.4.3.3. IBM Z network connectivity requirements	2566
18.4.3.4. Host machine resource requirements	2567
18.4.3.5. Minimum IBM Z system environment	2567
Hardware requirements	2567
Operating system requirements	2567
18.4.3.6. Minimum resource requirements	2567
18.4.3.7. Preferred IBM Z system environment	2568
Hardware requirements	2568
Operating system requirements	2568
18.4.3.8. Preferred resource requirements	2568
18.4.3.9. Certificate signing requests management	2569
18.4.3.10. Networking requirements for user-provisioned infrastructure	2569
18.4.3.10.1. Setting the cluster node hostnames through DHCP	2569
18.4.3.10.2. Network connectivity requirements	2570
NTP configuration for user-provisioned infrastructure	2571
18.4.3.11. User-provisioned DNS requirements	2571
18.4.3.11.1. Example DNS configuration for user-provisioned clusters	2573
18.4.3.12. Load balancing requirements for user-provisioned infrastructure	2575
18.4.3.12.1. Example load balancer configuration for user-provisioned clusters	2577
18.4.4. Preparing the user-provisioned infrastructure	2579
18.4.5. Validating DNS resolution for user-provisioned infrastructure	2581
18.4.6. Generating a key pair for cluster node SSH access	2583
18.4.7. Obtaining the installation program	2585
18.4.8. Installing the OpenShift CLI	2586
Installing the OpenShift CLI on Linux	2586
Installing the OpenShift CLI on Windows	2586
Installing the OpenShift CLI on macOS	2587
18.4.9. Manually creating the installation configuration file	2587
18.4.9.1. Sample install-config.yaml file for IBM Z	2588
18.4.9.2. Configuring the cluster-wide proxy during installation	2591
18.4.9.3. Configuring a three-node cluster	2593
18.4.10. Cluster Network Operator configuration	2594
18.4.10.1. Cluster Network Operator configuration object	2594
defaultNetwork object configuration	2595
Configuration for the OVN-Kubernetes network plugin	2595
kubeProxyConfig object configuration (OpenShiftSDN container network interface only)	2600
18.4.11. Creating the Kubernetes manifest and Ignition config files	2600
18.4.12. Installing RHCOS and starting the OpenShift Container Platform bootstrap process	2602
18.4.12.1. Installing RHCOS using IBM Secure Execution	2602
18.4.12.2. Configuring NBDE with static IP in an IBM Z or IBM LinuxONE environment	2605
18.4.12.3. Fast-track installation by using a prepackaged QCOW2 disk image	2607
18.4.12.4. Full installation on a new QCOW2 disk image	2608
18.4.12.5. Advanced RHCOS installation reference	2610
18.4.12.5.1. Networking options for ISO installations	2610
Configuring DHCP or static IP addresses	2610
Configuring an IP address without a static hostname	2611
Specifying multiple network interfaces	2611
Configuring default gateway and route	2611

Disabling DHCP on a single interface	2611
Combining DHCP and static IP configurations	2612
Configuring VLANs on individual interfaces	2612
Providing multiple DNS servers	2612
18.4.13. Waiting for the bootstrap process to complete	2612
18.4.14. Logging in to the cluster by using the CLI	2613
18.4.15. Approving the certificate signing requests for your machines	2614
18.4.16. Initial Operator configuration	2616
18.4.16.1. Image registry storage configuration	2617
18.4.16.1.1. Configuring registry storage for IBM Z	2617
18.4.16.1.2. Configuring storage for the image registry in non-production clusters	2619
18.4.17. Completing installation on user-provisioned infrastructure	2620
18.4.18. Telemetry access for OpenShift Container Platform	2622
18.4.19. Next steps	2622
18.5. INSTALLING A CLUSTER WITH RHEL KVM ON IBM Z AND IBM LINUXONE IN A RESTRICTED NETWORK	2622
18.5.1. Prerequisites	2623
18.5.2. About installations in restricted networks	2623
18.5.2.1. Additional limits	2624
18.5.3. Internet access for OpenShift Container Platform	2624
18.5.4. Machine requirements for a cluster with user-provisioned infrastructure	2624
18.5.4.1. Required machines	2624
18.5.4.2. Network connectivity requirements	2625
18.5.4.3. IBM Z network connectivity requirements	2625
18.5.4.4. Host machine resource requirements	2625
18.5.4.5. Minimum IBM Z system environment	2625
Hardware requirements	2626
Operating system requirements	2626
18.5.4.6. Minimum resource requirements	2626
18.5.4.7. Preferred IBM Z system environment	2626
Hardware requirements	2626
Operating system requirements	2627
18.5.4.8. Preferred resource requirements	2627
18.5.4.9. Certificate signing requests management	2627
18.5.4.10. Networking requirements for user-provisioned infrastructure	2627
18.5.4.10.1. Setting the cluster node hostnames through DHCP	2628
18.5.4.10.2. Network connectivity requirements	2628
NTP configuration for user-provisioned infrastructure	2629
18.5.4.11. User-provisioned DNS requirements	2630
18.5.4.11.1. Example DNS configuration for user-provisioned clusters	2631
18.5.4.12. Load balancing requirements for user-provisioned infrastructure	2634
18.5.4.12.1. Example load balancer configuration for user-provisioned clusters	2635
18.5.5. Preparing the user-provisioned infrastructure	2637
18.5.6. Validating DNS resolution for user-provisioned infrastructure	2639
18.5.7. Generating a key pair for cluster node SSH access	2642
18.5.8. Manually creating the installation configuration file	2643
18.5.8.1. Sample install-config.yaml file for IBM Z	2644
18.5.8.2. Configuring the cluster-wide proxy during installation	2647
18.5.8.3. Configuring a three-node cluster	2649
18.5.9. Cluster Network Operator configuration	2650
18.5.9.1. Cluster Network Operator configuration object	2651
defaultNetwork object configuration	2651
Configuration for the OVN-Kubernetes network plugin	2652

kubeProxyConfig object configuration (OpenShiftSDN container network interface only)	2656
18.5.10. Creating the Kubernetes manifest and Ignition config files	2657
18.5.11. Installing RHCOS and starting the OpenShift Container Platform bootstrap process	2659
18.5.11.1. Installing RHCOS using IBM Secure Execution	2659
18.5.11.2. Configuring NBDE with static IP in an IBM Z or IBM LinuxONE environment	2662
18.5.11.3. Fast-track installation by using a prepackaged QCOW2 disk image	2664
18.5.11.4. Full installation on a new QCOW2 disk image	2665
18.5.11.5. Advanced RHCOS installation reference	2667
18.5.11.5.1. Networking options for ISO installations	2667
Configuring DHCP or static IP addresses	2667
Configuring an IP address without a static hostname	2668
Specifying multiple network interfaces	2668
Configuring default gateway and route	2668
Disabling DHCP on a single interface	2669
Combining DHCP and static IP configurations	2669
Configuring VLANs on individual interfaces	2669
Providing multiple DNS servers	2669
18.5.12. Waiting for the bootstrap process to complete	2669
18.5.13. Logging in to the cluster by using the CLI	2670
18.5.14. Approving the certificate signing requests for your machines	2671
18.5.15. Initial Operator configuration	2673
18.5.15.1. Disabling the default OperatorHub catalog sources	2674
18.5.15.2. Image registry storage configuration	2675
18.5.15.2.1. Configuring registry storage for IBM Z	2675
18.5.15.2.2. Configuring storage for the image registry in non-production clusters	2677
18.5.16. Completing installation on user-provisioned infrastructure	2677
18.5.17. Next steps	2680
18.6. INSTALLING A CLUSTER IN AN LPAR ON IBM Z AND IBM LINUXONE	2680
18.6.1. Prerequisites	2680
18.6.2. Internet access for OpenShift Container Platform	2681
18.6.3. Requirements for a cluster with user-provisioned infrastructure	2681
18.6.3.1. Required machines for cluster installation	2681
18.6.3.2. Minimum resource requirements for cluster installation	2682
18.6.3.3. Minimum IBM Z system environment	2683
Hardware requirements	2683
Operating system requirements	2683
IBM Z network connectivity requirements	2683
Disk storage	2683
Storage / Main Memory	2684
18.6.3.4. Preferred IBM Z system environment	2684
Hardware requirements	2684
Operating system requirements	2684
IBM Z network connectivity requirements	2684
Disk storage	2684
Storage / Main Memory	2685
18.6.3.5. Certificate signing requests management	2685
18.6.3.6. Networking requirements for user-provisioned infrastructure	2685
18.6.3.6.1. Network connectivity requirements	2685
NTP configuration for user-provisioned infrastructure	2686
18.6.3.7. User-provisioned DNS requirements	2687
18.6.3.7.1. Example DNS configuration for user-provisioned clusters	2688
18.6.3.8. Load balancing requirements for user-provisioned infrastructure	2691
18.6.3.8.1. Example load balancer configuration for user-provisioned clusters	2692

18.6.4. Preparing the user-provisioned infrastructure	2694
18.6.5. Validating DNS resolution for user-provisioned infrastructure	2696
18.6.6. Generating a key pair for cluster node SSH access	2698
18.6.7. Obtaining the installation program	2700
18.6.8. Installing the OpenShift CLI	2700
Installing the OpenShift CLI on Linux	2701
Installing the OpenShift CLI on Windows	2701
Installing the OpenShift CLI on macOS	2702
18.6.9. Manually creating the installation configuration file	2702
18.6.9.1. Sample install-config.yaml file for IBM Z	2703
18.6.9.2. Configuring the cluster-wide proxy during installation	2706
18.6.9.3. Configuring a three-node cluster	2707
18.6.10. Cluster Network Operator configuration	2708
18.6.10.1. Cluster Network Operator configuration object	2709
defaultNetwork object configuration	2709
Configuration for the OVN-Kubernetes network plugin	2710
kubeProxyConfig object configuration (OpenShiftSDN container network interface only)	2714
18.6.11. Creating the Kubernetes manifest and Ignition config files	2715
18.6.12. Configuring NBDE with static IP in an IBM Z or IBM LinuxONE environment	2717
18.6.13. Installing RHCOS and starting the OpenShift Container Platform bootstrap process	2719
18.6.13.1. Advanced RHCOS installation reference	2722
18.6.13.1.1. Networking and bonding options for ISO installations	2722
Configuring DHCP or static IP addresses	2722
Configuring an IP address without a static hostname	2723
Specifying multiple network interfaces	2723
Configuring default gateway and route	2723
Disabling DHCP on a single interface	2723
Combining DHCP and static IP configurations	2724
Configuring VLANs on individual interfaces	2724
Providing multiple DNS servers	2724
Bonding multiple network interfaces to a single interface	2724
Bonding multiple network interfaces to a single interface	2725
Using network teaming	2725
18.6.14. Waiting for the bootstrap process to complete	2725
18.6.15. Logging in to the cluster by using the CLI	2726
18.6.16. Approving the certificate signing requests for your machines	2727
18.6.17. Initial Operator configuration	2729
18.6.17.1. Image registry storage configuration	2730
18.6.17.1.1. Configuring registry storage for IBM Z	2731
18.6.17.1.2. Configuring storage for the image registry in non-production clusters	2732
18.6.18. Completing installation on user-provisioned infrastructure	2733
18.6.19. Telemetry access for OpenShift Container Platform	2736
18.6.20. Next steps	2736
18.7. INSTALLING A CLUSTER IN AN LPAR ON IBM Z AND IBM LINUXONE IN A RESTRICTED NETWORK	2737
18.7.1. Prerequisites	2737
18.7.2. About installations in restricted networks	2737
18.7.2.1. Additional limits	2738
18.7.3. Internet access for OpenShift Container Platform	2738
18.7.4. Requirements for a cluster with user-provisioned infrastructure	2738
18.7.4.1. Required machines for cluster installation	2738
18.7.4.2. Minimum resource requirements for cluster installation	2739
18.7.4.3. Minimum IBM Z system environment	2740
Hardware requirements	2740

Operating system requirements	2741
IBM Z network connectivity requirements	2741
Disk storage	2741
Storage / Main Memory	2741
18.7.4.4. Preferred IBM Z system environment	2741
Hardware requirements	2741
Operating system requirements	2742
IBM Z network connectivity requirements	2742
Disk storage	2742
Storage / Main Memory	2742
18.7.4.5. Certificate signing requests management	2742
18.7.4.6. Networking requirements for user-provisioned infrastructure	2742
18.7.4.6.1. Setting the cluster node hostnames through DHCP	2743
18.7.4.6.2. Network connectivity requirements	2743
NTP configuration for user-provisioned infrastructure	2744
18.7.4.7. User-provisioned DNS requirements	2745
18.7.4.7.1. Example DNS configuration for user-provisioned clusters	2746
18.7.4.8. Load balancing requirements for user-provisioned infrastructure	2749
18.7.4.8.1. Example load balancer configuration for user-provisioned clusters	2750
18.7.5. Preparing the user-provisioned infrastructure	2752
18.7.6. Validating DNS resolution for user-provisioned infrastructure	2754
18.7.7. Generating a key pair for cluster node SSH access	2756
18.7.8. Manually creating the installation configuration file	2758
18.7.8.1. Sample install-config.yaml file for IBM Z	2758
18.7.8.2. Configuring the cluster-wide proxy during installation	2762
18.7.8.3. Configuring a three-node cluster	2763
18.7.9. Cluster Network Operator configuration	2764
18.7.9.1. Cluster Network Operator configuration object	2765
defaultNetwork object configuration	2765
Configuration for the OVN-Kubernetes network plugin	2766
kubeProxyConfig object configuration (OpenShiftSDN container network interface only)	2770
18.7.10. Creating the Kubernetes manifest and Ignition config files	2771
18.7.11. Configuring NBDE with static IP in an IBM Z or IBM LinuxONE environment	2773
18.7.12. Installing RHCOS and starting the OpenShift Container Platform bootstrap process	2775
18.7.12.1. Advanced RHCOS installation reference	2778
18.7.12.1.1. Networking and bonding options for ISO installations	2778
Configuring DHCP or static IP addresses	2778
Configuring an IP address without a static hostname	2779
Specifying multiple network interfaces	2779
Configuring default gateway and route	2779
Disabling DHCP on a single interface	2780
Combining DHCP and static IP configurations	2780
Configuring VLANs on individual interfaces	2780
Providing multiple DNS servers	2780
Bonding multiple network interfaces to a single interface	2780
Bonding multiple network interfaces to a single interface	2781
Using network teaming	2781
18.7.13. Waiting for the bootstrap process to complete	2781
18.7.14. Logging in to the cluster by using the CLI	2782
18.7.15. Approving the certificate signing requests for your machines	2783
18.7.16. Initial Operator configuration	2786
18.7.16.1. Disabling the default OperatorHub catalog sources	2786
18.7.16.2. Image registry storage configuration	2787

18.7.16.2.1. Configuring registry storage for IBM Z	2787
18.7.16.2.2. Configuring storage for the image registry in non-production clusters	2789
18.7.17. Completing installation on user-provisioned infrastructure	2789
18.7.18. Next steps	2793
18.8. INSTALLATION CONFIGURATION PARAMETERS FOR IBM Z AND IBM LINUXONE	2793
18.8.1. Available installation configuration parameters for IBM Z	2793
18.8.1.1. Required configuration parameters	2793
18.8.1.2. Network configuration parameters	2795
18.8.1.3. Optional configuration parameters	2797
CHAPTER 19. INSTALLING ON IBM POWER	2803
19.1. PREPARING TO INSTALL ON IBM POWER	2803
19.1.1. Prerequisites	2803
19.1.2. Choosing a method to install OpenShift Container Platform on IBM Power	2803
19.2. INSTALLING A CLUSTER ON IBM POWER	2803
19.2.1. Prerequisites	2803
19.2.2. Internet access for OpenShift Container Platform	2804
19.2.3. Requirements for a cluster with user-provisioned infrastructure	2804
19.2.3.1. Required machines for cluster installation	2804
19.2.3.2. Minimum resource requirements for cluster installation	2805
19.2.3.3. Minimum IBM Power requirements	2806
Hardware requirements	2806
Operating system requirements	2806
Disk storage for the IBM Power guest virtual machines	2806
Network for the PowerVM guest virtual machines	2807
Storage / main memory	2807
19.2.3.4. Recommended IBM Power system requirements	2807
Hardware requirements	2807
Operating system requirements	2807
Disk storage for the IBM Power guest virtual machines	2807
Network for the PowerVM guest virtual machines	2807
Storage / main memory	2807
19.2.3.5. Certificate signing requests management	2807
19.2.3.6. Networking requirements for user-provisioned infrastructure	2808
19.2.3.6.1. Setting the cluster node hostnames through DHCP	2808
19.2.3.6.2. Network connectivity requirements	2808
NTP configuration for user-provisioned infrastructure	2810
19.2.3.7. User-provisioned DNS requirements	2810
19.2.3.7.1. Example DNS configuration for user-provisioned clusters	2812
19.2.3.8. Load balancing requirements for user-provisioned infrastructure	2814
19.2.3.8.1. Example load balancer configuration for user-provisioned clusters	2816
19.2.4. Preparing the user-provisioned infrastructure	2818
19.2.5. Validating DNS resolution for user-provisioned infrastructure	2820
19.2.6. Generating a key pair for cluster node SSH access	2822
19.2.7. Obtaining the installation program	2824
19.2.8. Installing the OpenShift CLI	2824
Installing the OpenShift CLI on Linux	2825
Installing the OpenShift CLI on Windows	2825
Installing the OpenShift CLI on macOS	2826
19.2.9. Manually creating the installation configuration file	2826
19.2.9.1. Sample install-config.yaml file for IBM Power	2827
19.2.9.2. Configuring the cluster-wide proxy during installation	2829
19.2.9.3. Configuring a three-node cluster	2831

19.2.10. Cluster Network Operator configuration	2832
19.2.10.1. Cluster Network Operator configuration object	2832
defaultNetwork object configuration	2833
Configuration for the OVN-Kubernetes network plugin	2834
kubeProxyConfig object configuration (OpenShiftSDN container network interface only)	2838
19.2.11. Creating the Kubernetes manifest and Ignition config files	2839
19.2.12. Installing RHCOS and starting the OpenShift Container Platform bootstrap process	2841
19.2.12.1. Installing RHCOS by using an ISO image	2841
19.2.12.1.1. Advanced RHCOS installation reference	2844
19.2.12.1.1.1. Networking and bonding options for ISO installations	2844
Configuring DHCP or static IP addresses	2845
Configuring an IP address without a static hostname	2845
Specifying multiple network interfaces	2846
Configuring default gateway and route	2846
Disabling DHCP on a single interface	2846
Combining DHCP and static IP configurations	2846
Configuring VLANs on individual interfaces	2846
Providing multiple DNS servers	2847
Bonding multiple network interfaces to a single interface	2847
Bonding multiple SR-IOV network interfaces to a dual port NIC interface	2847
19.2.12.2. Installing RHCOS by using PXE booting	2848
19.2.12.3. Enabling multipathing with kernel arguments on RHCOS	2852
19.2.13. Waiting for the bootstrap process to complete	2853
19.2.14. Logging in to the cluster by using the CLI	2854
19.2.15. Approving the certificate signing requests for your machines	2855
19.2.16. Initial Operator configuration	2857
19.2.16.1. Image registry storage configuration	2858
19.2.16.1.1. Configuring registry storage for IBM Power	2859
19.2.16.1.2. Configuring storage for the image registry in non-production clusters	2860
19.2.17. Completing installation on user-provisioned infrastructure	2861
19.2.18. Telemetry access for OpenShift Container Platform	2863
19.2.19. Next steps	2863
19.3. INSTALLING A CLUSTER ON IBM POWER IN A RESTRICTED NETWORK	2863
19.3.1. Prerequisites	2864
19.3.2. About installations in restricted networks	2864
19.3.2.1. Additional limits	2865
19.3.3. Internet access for OpenShift Container Platform	2865
19.3.4. Requirements for a cluster with user-provisioned infrastructure	2865
19.3.4.1. Required machines for cluster installation	2865
19.3.4.2. Minimum resource requirements for cluster installation	2866
19.3.4.3. Minimum IBM Power requirements	2867
Hardware requirements	2867
Operating system requirements	2867
Disk storage for the IBM Power guest virtual machines	2867
Network for the PowerVM guest virtual machines	2868
Storage / main memory	2868
19.3.4.4. Recommended IBM Power system requirements	2868
Hardware requirements	2868
Operating system requirements	2868
Disk storage for the IBM Power guest virtual machines	2868
Network for the PowerVM guest virtual machines	2868
Storage / main memory	2868
19.3.4.5. Certificate signing requests management	2868

19.3.4.6. Networking requirements for user-provisioned infrastructure	2869
19.3.4.6.1. Setting the cluster node hostnames through DHCP	2869
19.3.4.6.2. Network connectivity requirements	2869
NTP configuration for user-provisioned infrastructure	2870
19.3.4.7. User-provisioned DNS requirements	2871
19.3.4.7.1. Example DNS configuration for user-provisioned clusters	2873
19.3.4.8. Load balancing requirements for user-provisioned infrastructure	2875
19.3.4.8.1. Example load balancer configuration for user-provisioned clusters	2877
19.3.5. Preparing the user-provisioned infrastructure	2879
19.3.6. Validating DNS resolution for user-provisioned infrastructure	2881
19.3.7. Generating a key pair for cluster node SSH access	2883
19.3.8. Manually creating the installation configuration file	2885
19.3.8.1. Sample install-config.yaml file for IBM Power	2886
19.3.8.2. Configuring the cluster-wide proxy during installation	2888
19.3.8.3. Configuring a three-node cluster	2890
19.3.9. Cluster Network Operator configuration	2891
19.3.9.1. Cluster Network Operator configuration object	2891
defaultNetwork object configuration	2892
Configuration for the OVN-Kubernetes network plugin	2893
kubeProxyConfig object configuration (OpenShiftSDN container network interface only)	2897
19.3.10. Creating the Kubernetes manifest and Ignition config files	2898
19.3.11. Installing RHCOS and starting the OpenShift Container Platform bootstrap process	2900
19.3.11.1. Installing RHCOS by using an ISO image	2900
19.3.11.1.1. Advanced RHCOS installation reference	2903
19.3.11.1.1.1. Networking and bonding options for ISO installations	2903
Configuring DHCP or static IP addresses	2904
Configuring an IP address without a static hostname	2904
Specifying multiple network interfaces	2905
Configuring default gateway and route	2905
Disabling DHCP on a single interface	2905
Combining DHCP and static IP configurations	2905
Configuring VLANs on individual interfaces	2905
Providing multiple DNS servers	2906
Bonding multiple network interfaces to a single interface	2906
Bonding multiple SR-IOV network interfaces to a dual port NIC interface	2906
19.3.11.2. Installing RHCOS by using PXE booting	2907
19.3.11.3. Enabling multipathing with kernel arguments on RHCOS	2911
19.3.12. Waiting for the bootstrap process to complete	2912
19.3.13. Logging in to the cluster by using the CLI	2913
19.3.14. Approving the certificate signing requests for your machines	2914
19.3.15. Initial Operator configuration	2916
19.3.15.1. Disabling the default OperatorHub catalog sources	2917
19.3.15.2. Image registry storage configuration	2918
19.3.15.2.1. Changing the image registry's management state	2918
19.3.15.2.2. Configuring registry storage for IBM Power	2918
19.3.15.2.3. Configuring storage for the image registry in non-production clusters	2920
19.3.16. Completing installation on user-provisioned infrastructure	2920
19.3.17. Next steps	2923
19.4. INSTALLATION CONFIGURATION PARAMETERS FOR IBM POWER	2923
19.4.1. Available installation configuration parameters for IBM Power	2923
19.4.1.1. Required configuration parameters	2923
19.4.1.2. Network configuration parameters	2925
19.4.1.3. Optional configuration parameters	2927

CHAPTER 20. INSTALLING ON IBM POWER VIRTUAL SERVER	2934
20.1. PREPARING TO INSTALL ON IBM POWER VIRTUAL SERVER	2934
20.1.1. Prerequisites	2934
20.1.2. Requirements for installing OpenShift Container Platform on IBM Power Virtual Server	2934
20.1.3. Choosing a method to install OpenShift Container Platform on IBM Power Virtual Server	2934
20.1.3.1. Installing a cluster on installer-provisioned infrastructure	2934
20.1.4. Configuring the Cloud Credential Operator utility	2935
20.1.5. Next steps	2936
20.2. CONFIGURING AN IBM CLOUD ACCOUNT	2936
20.2.1. Prerequisites	2936
20.2.2. Quotas and limits on IBM Power Virtual Server	2937
Virtual Private Cloud	2937
Application load balancer	2937
Transit Gateways	2937
Dynamic Host Configuration Protocol Service	2937
Networking	2937
Virtual Server Instances	2937
20.2.3. Configuring DNS resolution	2938
20.2.4. Using IBM Cloud Internet Services for DNS resolution	2938
20.2.5. IBM Cloud IAM Policies and API Key	2939
20.2.5.1. Pre-requisite permissions	2939
20.2.5.2. Cluster-creation permissions	2940
20.2.5.3. Access policy assignment	2940
20.2.5.4. Creating an API key	2941
20.2.6. Supported IBM Power Virtual Server regions and zones	2941
20.2.7. Next steps	2942
20.3. CREATING AN IBM POWER VIRTUAL SERVER WORKSPACE	2942
20.3.1. Creating an IBM Power Virtual Server workspace	2942
20.3.2. Next steps	2943
20.4. INSTALLING A CLUSTER ON IBM POWER VIRTUAL SERVER WITH CUSTOMIZATIONS	2943
20.4.1. Prerequisites	2943
20.4.2. Internet access for OpenShift Container Platform	2943
20.4.3. Generating a key pair for cluster node SSH access	2943
20.4.4. Obtaining the installation program	2945
20.4.5. Exporting the API key	2946
20.4.6. Creating the installation configuration file	2946
20.4.6.1. Sample customized install-config.yaml file for IBM Power Virtual Server	2948
20.4.6.2. Configuring the cluster-wide proxy during installation	2950
20.4.7. Manually creating IAM	2951
20.4.8. Deploying the cluster	2954
20.4.9. Installing the OpenShift CLI	2955
Installing the OpenShift CLI on Linux	2955
Installing the OpenShift CLI on Windows	2956
Installing the OpenShift CLI on macOS	2956
20.4.10. Logging in to the cluster by using the CLI	2957
20.4.11. Telemetry access for OpenShift Container Platform	2957
20.4.12. Next steps	2958
20.5. INSTALLING A CLUSTER ON IBM POWER VIRTUAL SERVER INTO AN EXISTING VPC	2958
20.5.1. Prerequisites	2958
20.5.2. About using a custom VPC	2958
20.5.2.1. Requirements for using your VPC	2958
20.5.2.2. VPC validation	2959
20.5.2.3. Isolation between clusters	2959

20.5.3. Internet access for OpenShift Container Platform	2959
20.5.4. Generating a key pair for cluster node SSH access	2960
20.5.5. Obtaining the installation program	2961
20.5.6. Exporting the API key	2962
20.5.7. Creating the installation configuration file	2963
20.5.7.1. Minimum resource requirements for cluster installation	2964
20.5.7.2. Sample customized install-config.yaml file for IBM Power Virtual Server	2965
20.5.7.3. Configuring the cluster-wide proxy during installation	2967
20.5.8. Manually creating IAM	2969
20.5.9. Deploying the cluster	2971
20.5.10. Installing the OpenShift CLI	2972
Installing the OpenShift CLI on Linux	2973
Installing the OpenShift CLI on Windows	2973
Installing the OpenShift CLI on macOS	2974
20.5.11. Logging in to the cluster by using the CLI	2974
20.5.12. Telemetry access for OpenShift Container Platform	2975
20.5.13. Next steps	2975
20.6. INSTALLING A PRIVATE CLUSTER ON IBM POWER VIRTUAL SERVER	2975
20.6.1. Prerequisites	2975
20.6.2. Private clusters	2976
20.6.3. Private clusters in IBM Power Virtual Server	2976
20.6.3.1. Limitations	2977
20.6.4. Requirements for using your VPC	2977
20.6.4.1. VPC validation	2977
20.6.4.2. Isolation between clusters	2978
20.6.5. Internet access for OpenShift Container Platform	2978
20.6.6. Generating a key pair for cluster node SSH access	2978
20.6.7. Obtaining the installation program	2980
20.6.8. Exporting the API key	2980
20.6.9. Manually creating the installation configuration file	2981
20.6.9.1. Minimum resource requirements for cluster installation	2982
20.6.9.2. Sample customized install-config.yaml file for IBM Power Virtual Server	2983
20.6.9.3. Configuring the cluster-wide proxy during installation	2985
20.6.10. Manually creating IAM	2986
20.6.11. Deploying the cluster	2989
20.6.12. Installing the OpenShift CLI	2990
Installing the OpenShift CLI on Linux	2990
Installing the OpenShift CLI on Windows	2991
Installing the OpenShift CLI on macOS	2991
20.6.13. Logging in to the cluster by using the CLI	2992
20.6.14. Telemetry access for OpenShift Container Platform	2993
20.6.15. Next steps	2993
20.7. INSTALLING A CLUSTER ON IBM POWER VIRTUAL SERVER IN A RESTRICTED NETWORK	2993
20.7.1. Prerequisites	2993
20.7.2. About installations in restricted networks	2994
20.7.2.1. Additional limits	2994
20.7.3. About using a custom VPC	2994
20.7.3.1. Requirements for using your VPC	2994
20.7.3.2. VPC validation	2995
20.7.3.3. Isolation between clusters	2995
20.7.4. Internet access for OpenShift Container Platform	2995
20.7.5. Generating a key pair for cluster node SSH access	2995
20.7.6. Exporting the API key	2997

20.7.7. Creating the installation configuration file	2997
20.7.7.1. Minimum resource requirements for cluster installation	3000
20.7.7.2. Sample customized install-config.yaml file for IBM Power Virtual Server	3001
20.7.7.3. Configuring the cluster-wide proxy during installation	3004
20.7.8. Manually creating IAM	3005
20.7.9. Deploying the cluster	3008
20.7.10. Installing the OpenShift CLI	3009
Installing the OpenShift CLI on Linux	3009
Installing the OpenShift CLI on Windows	3010
Installing the OpenShift CLI on macOS	3010
20.7.11. Logging in to the cluster by using the CLI	3011
20.7.12. Disabling the default OperatorHub catalog sources	3011
20.7.13. Telemetry access for OpenShift Container Platform	3012
20.7.14. Next steps	3012
20.8. UNINSTALLING A CLUSTER ON IBM POWER VIRTUAL SERVER	3012
20.8.1. Removing a cluster that uses installer-provisioned infrastructure	3012
20.9. INSTALLATION CONFIGURATION PARAMETERS FOR IBM POWER VIRTUAL SERVER	3014
20.9.1. Available installation configuration parameters for IBM Power Virtual Server	3014
20.9.1.1. Required configuration parameters	3014
20.9.1.2. Network configuration parameters	3016
20.9.1.3. Optional configuration parameters	3018
CHAPTER 21. INSTALLING ON OPENSTACK	3025
21.1. PREPARING TO INSTALL ON OPENSTACK	3025
21.1.1. Prerequisites	3025
21.1.2. Choosing a method to install OpenShift Container Platform on OpenStack	3025
21.1.2.1. Installing a cluster on installer-provisioned infrastructure	3025
21.1.2.2. Installing a cluster on user-provisioned infrastructure	3025
21.1.3. Scanning RHOSP endpoints for legacy HTTPS certificates	3026
21.1.3.1. Scanning RHOSP endpoints for legacy HTTPS certificates manually	3028
21.2. PREPARING TO INSTALL A CLUSTER THAT USES SR-IOV OR OVS-DPDK ON OPENSTACK	3029
21.2.1. Requirements for clusters on RHOSP that use either SR-IOV or OVS-DPDK	3029
21.2.1.1. Requirements for clusters on RHOSP that use SR-IOV	3029
21.2.1.2. Requirements for clusters on RHOSP that use OVS-DPDK	3030
21.2.2. Preparing to install a cluster that uses SR-IOV	3030
21.2.2.1. Creating SR-IOV networks for compute machines	3030
21.2.3. Preparing to install a cluster that uses OVS-DPDK	3031
21.2.4. Next steps	3031
21.3. INSTALLING A CLUSTER ON OPENSTACK WITH CUSTOMIZATIONS	3032
21.3.1. Prerequisites	3032
21.3.2. Resource guidelines for installing OpenShift Container Platform on RHOSP	3032
21.3.2.1. Control plane machines	3033
21.3.2.2. Compute machines	3034
21.3.2.3. Bootstrap machine	3034
21.3.2.4. Load balancing requirements for user-provisioned infrastructure	3034
21.3.2.4.1. Example load balancer configuration for clusters that are deployed with user-managed load balancers	3036
21.3.3. Internet access for OpenShift Container Platform	3038
21.3.4. Enabling Swift on RHOSP	3038
21.3.5. Configuring an image registry with custom storage on clusters that run on RHOSP	3039
21.3.6. Verifying external network access	3041
21.3.7. Defining parameters for the installation program	3042
21.3.8. Setting OpenStack Cloud Controller Manager options	3044

21.3.9. Obtaining the installation program	3045
21.3.10. Creating the installation configuration file	3046
21.3.10.1. Configuring the cluster-wide proxy during installation	3047
21.3.10.2. Custom subnets in RHOSP deployments	3049
21.3.10.3. Deploying a cluster with bare metal machines	3050
21.3.10.4. Cluster deployment on RHOSP provider networks	3051
21.3.10.4.1. RHOSP provider network requirements for cluster installation	3052
21.3.10.4.2. Deploying a cluster that has a primary interface on a provider network	3053
21.3.10.5. Sample customized install-config.yaml file for RHOSP	3055
21.3.10.6. Optional: Configuring a cluster with dual-stack networking	3057
21.3.10.6.1. Deploying the dual-stack cluster	3057
21.3.10.7. Installation configuration for a cluster on OpenStack with a user-managed load balancer	3060
21.3.11. Generating a key pair for cluster node SSH access	3061
21.3.12. Enabling access to the environment	3062
21.3.12.1. Enabling access with floating IP addresses	3062
21.3.12.2. Completing installation without floating IP addresses	3064
21.3.13. Deploying the cluster	3064
21.3.14. Verifying cluster status	3066
21.3.15. Logging in to the cluster by using the CLI	3066
21.3.16. Telemetry access for OpenShift Container Platform	3067
21.3.17. Next steps	3067
21.4. INSTALLING A CLUSTER ON OPENSTACK ON YOUR OWN INFRASTRUCTURE	3068
21.4.1. Prerequisites	3068
21.4.2. Internet access for OpenShift Container Platform	3068
21.4.3. Resource guidelines for installing OpenShift Container Platform on RHOSP	3069
21.4.3.1. Control plane machines	3070
21.4.3.2. Compute machines	3070
21.4.3.3. Bootstrap machine	3070
21.4.4. Downloading playbook dependencies	3071
21.4.5. Downloading the installation playbooks	3071
21.4.6. Obtaining the installation program	3073
21.4.7. Generating a key pair for cluster node SSH access	3073
21.4.8. Creating the Red Hat Enterprise Linux CoreOS (RHCOS) image	3075
21.4.9. Verifying external network access	3076
21.4.10. Enabling access to the environment	3077
21.4.10.1. Enabling access with floating IP addresses	3077
21.4.10.2. Completing installation without floating IP addresses	3079
21.4.11. Defining parameters for the installation program	3079
21.4.12. Creating network resources on RHOSP	3080
21.4.13. Creating the installation configuration file	3081
21.4.13.1. Custom subnets in RHOSP deployments	3082
21.4.13.2. Sample customized install-config.yaml file for RHOSP	3083
21.4.13.3. Setting a custom subnet for machines	3085
21.4.13.4. Emptying compute machine pools	3086
21.4.13.5. Cluster deployment on RHOSP provider networks	3086
21.4.13.5.1. RHOSP provider network requirements for cluster installation	3088
21.4.13.5.2. Deploying a cluster that has a primary interface on a provider network	3089
21.4.14. Creating the Kubernetes manifest and Ignition config files	3090
21.4.15. Preparing the bootstrap Ignition files	3092
21.4.16. Creating control plane Ignition config files on RHOSP	3094
21.4.17. Updating network resources on RHOSP	3095
21.4.17.1. Deploying a cluster with bare metal machines	3097
21.4.18. Creating the bootstrap machine on RHOSP	3098

21.4.19. Creating the control plane machines on RHOSP	3099
21.4.20. Logging in to the cluster by using the CLI	3099
21.4.21. Deleting bootstrap resources from RHOSP	3100
21.4.22. Creating compute machines on RHOSP	3101
21.4.23. Approving the certificate signing requests for your machines	3101
21.4.24. Verifying a successful installation	3104
21.4.25. Telemetry access for OpenShift Container Platform	3104
21.4.26. Next steps	3105
21.5. INSTALLING A CLUSTER ON OPENSTACK IN A RESTRICTED NETWORK	3105
21.5.1. Prerequisites	3105
21.5.2. About installations in restricted networks	3105
21.5.2.1. Additional limits	3106
21.5.3. Resource guidelines for installing OpenShift Container Platform on RHOSP	3106
21.5.3.1. Control plane machines	3107
21.5.3.2. Compute machines	3107
21.5.3.3. Bootstrap machine	3108
21.5.4. Internet access for OpenShift Container Platform	3108
21.5.5. Enabling Swift on RHOSP	3108
21.5.6. Defining parameters for the installation program	3109
21.5.7. Setting OpenStack Cloud Controller Manager options	3110
21.5.8. Creating the RHCOS image for restricted network installations	3112
21.5.9. Creating the installation configuration file	3113
21.5.9.1. Configuring the cluster-wide proxy during installation	3116
21.5.9.2. Sample customized install-config.yaml file for restricted OpenStack installations	3117
21.5.10. Generating a key pair for cluster node SSH access	3118
21.5.11. Enabling access to the environment	3120
21.5.11.1. Enabling access with floating IP addresses	3120
21.5.11.2. Completing installation without floating IP addresses	3121
21.5.12. Deploying the cluster	3122
21.5.13. Verifying cluster status	3123
21.5.14. Logging in to the cluster by using the CLI	3124
21.5.15. Disabling the default OperatorHub catalog sources	3125
21.5.16. Telemetry access for OpenShift Container Platform	3125
21.5.17. Next steps	3125
21.6. OPENSTACK CLOUD CONTROLLER MANAGER REFERENCE GUIDE	3126
21.6.1. The OpenStack Cloud Controller Manager	3126
21.6.2. The OpenStack Cloud Controller Manager (CCM) config map	3126
21.6.2.1. Load balancer options	3127
21.6.2.2. Options that the Operator overrides	3130
21.7. DEPLOYING ON OPENSTACK WITH ROOTVOLUME AND ETCD ON LOCAL DISK	3132
21.7.1. Deploying RHOSP on local disk	3132
21.7.2. Additional resources	3138
21.8. UNINSTALLING A CLUSTER ON OPENSTACK	3138
21.8.1. Removing a cluster that uses installer-provisioned infrastructure	3138
21.9. UNINSTALLING A CLUSTER ON RHOSP FROM YOUR OWN INFRASTRUCTURE	3139
21.9.1. Downloading playbook dependencies	3139
21.9.2. Removing a cluster from RHOSP that uses your own infrastructure	3140
21.10. INSTALLATION CONFIGURATION PARAMETERS FOR OPENSTACK	3140
21.10.1. Available installation configuration parameters for OpenStack	3140
21.10.1.1. Required configuration parameters	3141
21.10.1.2. Network configuration parameters	3142
21.10.1.3. Optional configuration parameters	3144
21.10.1.4. Additional Red Hat OpenStack Platform (RHOSP) configuration parameters	3150

21.10.1.5. Optional RHOSP configuration parameters	3153
CHAPTER 22. INSTALLING ON OCI	3158
22.1. INSTALLING A CLUSTER ON ORACLE CLOUD INFRASTRUCTURE (OCI) BY USING THE ASSISTED INSTALLER	3158
22.1.1. The Assisted Installer and OCI overview	3158
22.1.2. Creating OCI resources and services	3159
22.1.3. Using the Assisted Installer to generate an OCI-compatible discovery ISO image	3159
22.1.4. Provisioning OCI infrastructure for your cluster	3161
22.1.5. Completing the remaining Assisted Installer steps	3161
22.1.6. Verifying a successful cluster installation on OCI	3162
22.1.7. Troubleshooting the installation of a cluster on OCI	3163
The Ingress Load Balancer in OCI is not at a healthy status	3163
OCI create stack operation fails with an Error: 400-InvalidParameter message	3163
22.2. INSTALLING A CLUSTER ON ORACLE CLOUD INFRASTRUCTURE (OCI) BY USING THE AGENT-BASED INSTALLER	3164
22.2.1. The Agent-based Installer and OCI overview	3164
22.2.2. Creating OCI infrastructure resources and services	3166
22.2.3. Creating configuration files for installing a cluster on OCI	3167
22.2.4. Configuring your firewall for OpenShift Container Platform	3171
22.2.5. Running a cluster on OCI	3174
22.2.6. Verifying that your Agent-based cluster installation runs on OCI	3174
CHAPTER 23. INSTALLING ON VSPHERE	3177
23.1. INSTALLATION METHODS	3177
23.1.1. Assisted Installer	3177
23.1.2. Agent-based Installer	3177
23.1.3. Installer-provisioned infrastructure installation	3177
23.1.4. User-provisioned infrastructure installation	3177
23.1.5. Additional resources	3178
23.2. INSTALLER-PROVISIONED INFRASTRUCTURE	3178
23.2.1. vSphere installation requirements	3178
23.2.1.1. VMware vSphere infrastructure requirements	3178
23.2.1.2. Network connectivity requirements	3180
23.2.1.3. VMware vSphere CSI Driver Operator requirements	3181
23.2.1.4. vCenter requirements	3182
Required vCenter account privileges	3182
Minimum required vCenter account privileges	3190
Using OpenShift Container Platform with vMotion	3199
Cluster resources	3200
Cluster limits	3200
Networking requirements	3201
Required IP Addresses	3201
DNS records	3201
Static IP addresses for vSphere nodes	3202
23.2.2. Preparing to install a cluster using installer-provisioned infrastructure	3203
23.2.2.1. Obtaining the installation program	3203
23.2.2.2. Installing the OpenShift CLI	3204
Installing the OpenShift CLI on Linux	3204
Installing the OpenShift CLI on Windows	3205
Installing the OpenShift CLI on macOS	3205
23.2.2.3. Generating a key pair for cluster node SSH access	3206
23.2.2.4. Adding vCenter root CA certificates to your system trust	3208

23.2.3. Installing a cluster on vSphere	3209
23.2.3.1. Prerequisites	3209
23.2.3.2. Internet access for OpenShift Container Platform	3209
23.2.3.3. Deploying the cluster	3210
23.2.3.4. Logging in to the cluster by using the CLI	3213
23.2.3.5. Creating registry storage	3213
23.2.3.5.1. Image registry removed during installation	3213
23.2.3.5.2. Image registry storage configuration	3214
23.2.3.5.2.1. Configuring registry storage for VMware vSphere	3214
23.2.3.5.2.2. Configuring block registry storage for VMware vSphere	3215
23.2.3.6. Telemetry access for OpenShift Container Platform	3217
23.2.3.7. Next steps	3217
23.2.4. Installing a cluster on vSphere with customizations	3217
23.2.4.1. Prerequisites	3217
23.2.4.2. Internet access for OpenShift Container Platform	3218
23.2.4.3. VMware vSphere region and zone enablement	3218
23.2.4.4. Creating the installation configuration file	3220
23.2.4.4.1. Sample install-config.yaml file for an installer-provisioned VMware vSphere cluster	3222
23.2.4.4.2. Configuring the cluster-wide proxy during installation	3224
23.2.4.4.3. Configuring regions and zones for a VMware vCenter	3226
23.2.4.5. Services for a user-managed load balancer	3228
23.2.4.5.1. Configuring a user-managed load balancer	3231
23.2.4.6. Deploying the cluster	3238
23.2.4.7. Logging in to the cluster by using the CLI	3239
23.2.4.8. Creating registry storage	3240
23.2.4.8.1. Image registry removed during installation	3240
23.2.4.8.2. Image registry storage configuration	3240
23.2.4.8.2.1. Configuring registry storage for VMware vSphere	3241
23.2.4.8.2.2. Configuring block registry storage for VMware vSphere	3242
23.2.4.9. Telemetry access for OpenShift Container Platform	3243
23.2.4.10. Next steps	3244
23.2.5. Installing a cluster on vSphere with network customizations	3244
23.2.5.1. Prerequisites	3244
23.2.5.2. Internet access for OpenShift Container Platform	3245
23.2.5.3. VMware vSphere region and zone enablement	3245
23.2.5.4. Creating the installation configuration file	3247
23.2.5.4.1. Sample install-config.yaml file for an installer-provisioned VMware vSphere cluster	3249
23.2.5.4.2. Configuring the cluster-wide proxy during installation	3251
23.2.5.4.3. Optional: Deploying with dual-stack networking	3253
23.2.5.4.4. Configuring regions and zones for a VMware vCenter	3253
23.2.5.5. Network configuration phases	3256
23.2.5.6. Specifying advanced network configuration	3256
23.2.5.7. Cluster Network Operator configuration	3257
23.2.5.7.1. Cluster Network Operator configuration object	3258
defaultNetwork object configuration	3259
Configuration for the OVN-Kubernetes network plugin	3259
kubeProxyConfig object configuration (OpenShiftSDN container network interface only)	3264
23.2.5.8. Services for a user-managed load balancer	3264
23.2.5.8.1. Configuring a user-managed load balancer	3267
23.2.5.9. Deploying the cluster	3274
23.2.5.10. Logging in to the cluster by using the CLI	3276
23.2.5.11. Creating registry storage	3277
23.2.5.11.1. Image registry removed during installation	3277

23.2.5.11.2. Image registry storage configuration	3277
23.2.5.11.2.1. Configuring registry storage for VMware vSphere	3277
23.2.5.11.2.2. Configuring block registry storage for VMware vSphere	3279
23.2.5.12. Telemetry access for OpenShift Container Platform	3280
23.2.5.13. Configuring network components to run on the control plane	3280
23.2.5.14. Next steps	3282
23.2.6. Installing a cluster on vSphere in a restricted network	3283
23.2.6.1. Prerequisites	3283
23.2.6.2. About installations in restricted networks	3283
23.2.6.2.1. Additional limits	3284
23.2.6.3. Internet access for OpenShift Container Platform	3284
23.2.6.4. Creating the RHCOS image for restricted network installations	3284
23.2.6.5. VMware vSphere region and zone enablement	3285
23.2.6.6. Creating the installation configuration file	3286
23.2.6.6.1. Sample install-config.yaml file for an installer-provisioned VMware vSphere cluster	3290
23.2.6.6.2. Configuring the cluster-wide proxy during installation	3292
23.2.6.6.3. Configuring regions and zones for a VMware vCenter	3294
23.2.6.7. Services for a user-managed load balancer	3296
23.2.6.7.1. Configuring a user-managed load balancer	3299
23.2.6.8. Deploying the cluster	3306
23.2.6.9. Logging in to the cluster by using the CLI	3307
23.2.6.10. Disabling the default OperatorHub catalog sources	3308
23.2.6.11. Creating registry storage	3308
23.2.6.11.1. Image registry removed during installation	3309
23.2.6.11.2. Image registry storage configuration	3309
23.2.6.11.2.1. Configuring registry storage for VMware vSphere	3309
23.2.6.12. Telemetry access for OpenShift Container Platform	3310
23.2.6.13. Next steps	3311
23.3. USER-PROVISIONED INFRASTRUCTURE	3311
23.3.1. vSphere installation requirements for user-provisioned infrastructure	3311
23.3.1.1. VMware vSphere infrastructure requirements	3311
23.3.1.2. VMware vSphere CSI Driver Operator requirements	3313
23.3.1.3. Requirements for a cluster with user-provisioned infrastructure	3314
23.3.1.3.1. vCenter requirements	3314
Required vCenter account privileges	3314
Minimum required vCenter account privileges	3322
Using OpenShift Container Platform with vMotion	3327
Cluster resources	3328
Cluster limits	3328
Networking requirements	3329
DNS records	3329
23.3.1.3.2. Required machines for cluster installation	3330
23.3.1.3.3. Minimum resource requirements for cluster installation	3331
23.3.1.3.4. Requirements for encrypting virtual machines	3332
23.3.1.3.5. Certificate signing requests management	3332
23.3.1.3.6. Networking requirements for user-provisioned infrastructure	3333
23.3.1.3.6.1. Setting the cluster node hostnames through DHCP	3333
23.3.1.3.6.2. Network connectivity requirements	3333
Ethernet adaptor hardware address requirements	3335
NTP configuration for user-provisioned infrastructure	3335
23.3.1.3.7. User-provisioned DNS requirements	3335
23.3.1.3.7.1. Example DNS configuration for user-provisioned clusters	3337
23.3.1.3.8. Load balancing requirements for user-provisioned infrastructure	3339

23.3.1.3.8.1. Example load balancer configuration for user-provisioned clusters	3341
23.3.2. Preparing to install a cluster using user-provisioned infrastructure	3343
23.3.2.1. Obtaining the installation program	3343
23.3.2.2. Installing the OpenShift CLI	3344
Installing the OpenShift CLI on Linux	3344
Installing the OpenShift CLI on Windows	3345
Installing the OpenShift CLI on macOS	3345
23.3.2.3. Generating a key pair for cluster node SSH access	3346
23.3.2.4. Preparing the user-provisioned infrastructure	3348
23.3.2.5. Validating DNS resolution for user-provisioned infrastructure	3350
23.3.3. Installing a cluster on vSphere with user-provisioned infrastructure	3352
23.3.3.1. Prerequisites	3352
23.3.3.2. Internet access for OpenShift Container Platform	3353
23.3.3.3. VMware vSphere region and zone enablement	3353
23.3.3.4. Manually creating the installation configuration file	3355
23.3.3.4.1. Sample install-config.yaml file for VMware vSphere	3356
23.3.3.4.2. Configuring the cluster-wide proxy during installation	3358
23.3.3.4.3. Configuring regions and zones for a VMware vCenter	3360
23.3.3.5. Creating the Kubernetes manifest and Ignition config files	3362
23.3.3.6. Extracting the infrastructure name	3364
23.3.3.7. Installing RHCOS and starting the OpenShift Container Platform bootstrap process	3365
23.3.3.8. Adding more compute machines to a cluster in vSphere	3370
23.3.3.9. Disk partitioning	3371
Creating a separate /var partition	3371
23.3.3.10. Waiting for the bootstrap process to complete	3373
23.3.3.11. Logging in to the cluster by using the CLI	3374
23.3.3.12. Approving the certificate signing requests for your machines	3375
23.3.3.13. Initial Operator configuration	3378
23.3.3.13.1. Image registry removed during installation	3378
23.3.3.13.2. Image registry storage configuration	3379
23.3.3.13.2.1. Configuring registry storage for VMware vSphere	3379
23.3.3.13.2.2. Configuring storage for the image registry in non-production clusters	3380
23.3.3.13.2.3. Configuring block registry storage for VMware vSphere	3381
23.3.3.14. Completing installation on user-provisioned infrastructure	3382
23.3.3.15. Configuring vSphere DRS anti-affinity rules for control plane nodes	3385
23.3.3.16. Telemetry access for OpenShift Container Platform	3386
23.3.3.17. Next steps	3386
23.3.4. Installing a cluster on vSphere with network customizations	3386
23.3.4.1. Prerequisites	3387
23.3.4.2. Internet access for OpenShift Container Platform	3387
23.3.4.3. VMware vSphere region and zone enablement	3388
23.3.4.4. Manually creating the installation configuration file	3389
23.3.4.4.1. Sample install-config.yaml file for VMware vSphere	3390
23.3.4.4.2. Configuring the cluster-wide proxy during installation	3393
23.3.4.4.3. Configuring regions and zones for a VMware vCenter	3395
23.3.4.5. Network configuration phases	3397
23.3.4.6. Specifying advanced network configuration	3398
23.3.4.7. Cluster Network Operator configuration	3399
23.3.4.7.1. Cluster Network Operator configuration object	3399
defaultNetwork object configuration	3400
Configuration for the OVN-Kubernetes network plugin	3400
kubeProxyConfig object configuration (OpenShiftSDN container network interface only)	3405
23.3.4.8. Creating the Ignition config files	3405

23.3.4.9. Extracting the infrastructure name	3406
23.3.4.10. Installing RHCOS and starting the OpenShift Container Platform bootstrap process	3407
23.3.4.11. Adding more compute machines to a cluster in vSphere	3412
23.3.4.12. Disk partitioning	3413
Creating a separate /var partition	3414
23.3.4.13. Waiting for the bootstrap process to complete	3416
23.3.4.14. Logging in to the cluster by using the CLI	3417
23.3.4.15. Approving the certificate signing requests for your machines	3417
23.3.4.15.1. Initial Operator configuration	3420
23.3.4.15.2. Image registry removed during installation	3421
23.3.4.15.3. Image registry storage configuration	3421
23.3.4.15.3.1. Configuring block registry storage for VMware vSphere	3421
23.3.4.16. Completing installation on user-provisioned infrastructure	3422
23.3.4.17. Configuring vSphere DRS anti-affinity rules for control plane nodes	3425
23.3.4.18. Telemetry access for OpenShift Container Platform	3426
23.3.4.19. Next steps	3426
23.3.5. Installing a cluster on vSphere in a restricted network with user-provisioned infrastructure	3426
23.3.5.1. Prerequisites	3427
23.3.5.2. About installations in restricted networks	3428
23.3.5.2.1. Additional limits	3428
23.3.5.3. Internet access for OpenShift Container Platform	3428
23.3.5.4. VMware vSphere region and zone enablement	3428
23.3.5.5. Manually creating the installation configuration file	3430
23.3.5.5.1. Sample install-config.yaml file for VMware vSphere	3432
23.3.5.5.2. Configuring the cluster-wide proxy during installation	3435
23.3.5.5.3. Configuring regions and zones for a VMware vCenter	3436
23.3.5.6. Creating the Kubernetes manifest and Ignition config files	3439
23.3.5.7. Configuring chrony time service	3440
23.3.5.8. Extracting the infrastructure name	3441
23.3.5.9. Installing RHCOS and starting the OpenShift Container Platform bootstrap process	3442
23.3.5.10. Adding more compute machines to a cluster in vSphere	3447
23.3.5.11. Disk partitioning	3448
Creating a separate /var partition	3448
23.3.5.12. Waiting for the bootstrap process to complete	3450
23.3.5.13. Logging in to the cluster by using the CLI	3451
23.3.5.14. Approving the certificate signing requests for your machines	3452
23.3.5.15. Initial Operator configuration	3454
23.3.5.15.1. Disabling the default OperatorHub catalog sources	3455
23.3.5.15.2. Image registry storage configuration	3456
23.3.5.15.2.1. Configuring registry storage for VMware vSphere	3456
23.3.5.15.2.2. Configuring storage for the image registry in non-production clusters	3457
23.3.5.15.2.3. Configuring block registry storage for VMware vSphere	3458
23.3.5.16. Completing installation on user-provisioned infrastructure	3459
23.3.5.17. Configuring vSphere DRS anti-affinity rules for control plane nodes	3462
23.3.5.18. Telemetry access for OpenShift Container Platform	3463
23.3.5.19. Next steps	3463
23.4. INSTALLING A CLUSTER ON VSPHERE USING THE ASSISTED INSTALLER	3463
23.4.1. Additional resources	3464
23.5. INSTALLING A CLUSTER ON VSPHERE USING THE AGENT-BASED INSTALLER	3464
23.5.1. Additional resources	3464
23.6. INSTALLING A THREE-NODE CLUSTER ON VSPHERE	3464
23.6.1. Configuring a three-node cluster	3464
23.6.2. Next steps	3465

23.7. UNINSTALLING A CLUSTER ON VSPHERE THAT USES INSTALLER-PROVISIONED INFRASTRUCTURE	3465
23.7.1. Removing a cluster that uses installer-provisioned infrastructure	3465
23.8. USING THE VSPHERE PROBLEM DETECTOR OPERATOR	3466
23.8.1. About the vSphere Problem Detector Operator	3466
23.8.2. Running the vSphere Problem Detector Operator checks	3467
23.8.3. Viewing the events from the vSphere Problem Detector Operator	3468
23.8.4. Viewing the logs from the vSphere Problem Detector Operator	3468
23.8.5. Configuration checks run by the vSphere Problem Detector Operator	3469
23.8.6. About the storage class configuration check	3470
23.8.7. Metrics for the vSphere Problem Detector Operator	3471
23.8.8. Additional resources	3471
23.9. INSTALLATION CONFIGURATION PARAMETERS FOR VSPHERE	3471
23.9.1. Available installation configuration parameters for vSphere	3471
23.9.1.1. Required configuration parameters	3472
23.9.1.2. Network configuration parameters	3473
23.9.1.3. Optional configuration parameters	3475
23.9.1.4. Additional VMware vSphere configuration parameters	3481
23.9.1.5. Deprecated VMware vSphere configuration parameters	3484
23.9.1.6. Optional VMware vSphere machine pool configuration parameters	3486
CHAPTER 24. INSTALLING ON ANY PLATFORM	3488
24.1. INSTALLING A CLUSTER ON ANY PLATFORM	3488
24.1.1. Prerequisites	3488
24.1.2. Internet access for OpenShift Container Platform	3488
24.1.3. Requirements for a cluster with user-provisioned infrastructure	3488
24.1.3.1. Required machines for cluster installation	3489
24.1.3.2. Minimum resource requirements for cluster installation	3489
24.1.3.3. Certificate signing requests management	3490
24.1.3.4. Networking requirements for user-provisioned infrastructure	3490
24.1.3.4.1. Setting the cluster node hostnames through DHCP	3491
24.1.3.4.2. Network connectivity requirements	3491
NTP configuration for user-provisioned infrastructure	3492
24.1.3.5. User-provisioned DNS requirements	3493
24.1.3.5.1. Example DNS configuration for user-provisioned clusters	3495
24.1.3.6. Load balancing requirements for user-provisioned infrastructure	3497
24.1.3.6.1. Example load balancer configuration for user-provisioned clusters	3499
24.1.4. Preparing the user-provisioned infrastructure	3500
24.1.5. Validating DNS resolution for user-provisioned infrastructure	3502
24.1.6. Generating a key pair for cluster node SSH access	3505
24.1.7. Obtaining the installation program	3506
24.1.8. Installing the OpenShift CLI	3507
Installing the OpenShift CLI on Linux	3507
Installing the OpenShift CLI on Windows	3508
Installing the OpenShift CLI on macOS	3508
24.1.9. Manually creating the installation configuration file	3509
24.1.9.1. Sample install-config.yaml file for other platforms	3510
24.1.9.2. Configuring the cluster-wide proxy during installation	3512
24.1.9.3. Configuring a three-node cluster	3514
24.1.10. Creating the Kubernetes manifest and Ignition config files	3515
24.1.11. Installing RHCOS and starting the OpenShift Container Platform bootstrap process	3517
24.1.11.1. Installing RHCOS by using an ISO image	3518
24.1.11.2. Installing RHCOS by using PXE or iPXE booting	3521

24.1.11.3. Advanced RHCOS installation configuration	3526
24.1.11.3.1. Using advanced networking options for PXE and ISO installations	3526
24.1.11.3.2. Disk partitioning	3527
24.1.11.3.2.1. Creating a separate /var partition	3528
24.1.11.3.2.2. Retaining existing partitions	3530
24.1.11.3.3. Identifying Ignition configs	3531
24.1.11.3.4. Advanced RHCOS installation reference	3531
24.1.11.3.4.1. Networking and bonding options for ISO installations	3531
Configuring DHCP or static IP addresses	3532
Configuring an IP address without a static hostname	3532
Specifying multiple network interfaces	3533
Configuring default gateway and route	3533
Disabling DHCP on a single interface	3533
Combining DHCP and static IP configurations	3533
Configuring VLANs on individual interfaces	3533
Providing multiple DNS servers	3534
Bonding multiple network interfaces to a single interface	3534
Bonding multiple SR-IOV network interfaces to a dual port NIC interface	3534
Using network teaming	3535
24.1.11.3.4.2. coreos-installer options for ISO and PXE installations	3535
24.1.11.3.4.3. coreos.inst boot options for ISO or PXE installations	3540
24.1.12. Waiting for the bootstrap process to complete	3541
24.1.13. Logging in to the cluster by using the CLI	3542
24.1.14. Approving the certificate signing requests for your machines	3543
24.1.15. Initial Operator configuration	3546
24.1.15.1. Disabling the default OperatorHub catalog sources	3547
24.1.15.2. Image registry removed during installation	3547
24.1.15.3. Image registry storage configuration	3547
24.1.15.3.1. Configuring registry storage for bare metal and other manual installations	3547
24.1.15.3.2. Configuring storage for the image registry in non-production clusters	3549
24.1.15.3.3. Configuring block registry storage for bare metal	3549
24.1.16. Completing installation on user-provisioned infrastructure	3551
24.1.17. Telemetry access for OpenShift Container Platform	3553
24.1.18. Next steps	3553
CHAPTER 25. INSTALLATION CONFIGURATION	3554
25.1. CUSTOMIZING NODES	3554
25.1.1. Creating machine configs with Butane	3554
25.1.1.1. About Butane	3554
25.1.1.2. Installing Butane	3554
25.1.1.3. Creating a MachineConfig object by using Butane	3555
25.1.2. Adding day-1 kernel arguments	3556
25.1.3. Adding kernel modules to nodes	3557
25.1.3.1. Building and testing the kernel module container	3558
25.1.3.2. Provisioning a kernel module to OpenShift Container Platform	3561
25.1.3.2.1. Provision kernel modules via a MachineConfig object	3561
25.1.4. Encrypting and mirroring disks during installation	3563
25.1.4.1. About disk encryption	3563
25.1.4.1.1. Configuring an encryption threshold	3564
25.1.4.2. About disk mirroring	3566
25.1.4.3. Configuring disk encryption and mirroring	3566
25.1.4.4. Configuring a RAID-enabled data volume	3573
25.1.4.5. Configuring an Intel® Virtual RAID on CPU (VROC) data volume	3575

25.1.5. Configuring chrony time service	3577
25.1.6. Additional resources	3578
25.2. CONFIGURING YOUR FIREWALL	3578
25.2.1. Configuring your firewall for OpenShift Container Platform	3578
25.2.2. OpenShift Container Platform network flow matrix	3583
25.3. ENABLING LINUX CONTROL GROUP VERSION 1 (CGROUP V1)	3591
25.3.1. Enabling Linux cgroup v1 during installation	3591
CHAPTER 26. VALIDATING AN INSTALLATION	3593
26.1. REVIEWING THE INSTALLATION LOG	3593
26.2. VIEWING THE IMAGE PULL SOURCE	3593
26.3. GETTING CLUSTER VERSION, STATUS, AND UPDATE DETAILS	3594
26.4. VERIFYING THAT A CLUSTER USES SHORT-TERM CREDENTIALS	3596
26.5. QUERYING THE STATUS OF THE CLUSTER NODES BY USING THE CLI	3597
26.6. REVIEWING THE CLUSTER STATUS FROM THE OPENSIFT CONTAINER PLATFORM WEB CONSOLE	3598
26.7. REVIEWING THE CLUSTER STATUS FROM RED HAT OPENSIFT CLUSTER MANAGER	3599
26.8. CHECKING CLUSTER RESOURCE AVAILABILITY AND UTILIZATION	3600
26.9. LISTING ALERTS THAT ARE FIRING	3601
26.10. NEXT STEPS	3601
CHAPTER 27. TROUBLESHOOTING INSTALLATION ISSUES	3603
27.1. PREREQUISITES	3603
27.2. GATHERING LOGS FROM A FAILED INSTALLATION	3603
27.3. MANUALLY GATHERING LOGS WITH SSH ACCESS TO YOUR HOST(S)	3604
27.4. MANUALLY GATHERING LOGS WITHOUT SSH ACCESS TO YOUR HOST(S)	3605
27.5. GETTING DEBUG INFORMATION FROM THE INSTALLATION PROGRAM	3605
27.6. REINSTALLING THE OPENSIFT CONTAINER PLATFORM CLUSTER	3606
CHAPTER 28. SUPPORT FOR FIPS CRYPTOGRAPHY	3607
28.1. OBTAINING A FIPS-CAPABLE INSTALLATION PROGRAM USING OC ADM EXTRACT	3607
28.2. OBTAINING A FIPS-CAPABLE INSTALLATION PROGRAM USING THE PUBLIC OPENSIFT MIRROR	3608
28.3. FIPS VALIDATION IN OPENSIFT CONTAINER PLATFORM	3608
28.4. FIPS SUPPORT IN COMPONENTS THAT THE CLUSTER USES	3609
28.4.1. etcd	3609
28.4.2. Storage	3609
28.4.3. Runtimes	3609
28.5. INSTALLING A CLUSTER IN FIPS MODE	3609

CHAPTER 1. OPENSIFT CONTAINER PLATFORM INSTALLATION OVERVIEW

1.1. ABOUT OPENSIFT CONTAINER PLATFORM INSTALLATION

The OpenShift Container Platform installation program offers four methods for deploying a cluster which are detailed in the following list:

- **Interactive:** You can deploy a cluster with the web-based [Assisted Installer](#). This is an ideal approach for clusters with networks connected to the internet. The Assisted Installer is the easiest way to install OpenShift Container Platform, it provides smart defaults, and it performs pre-flight validations before installing the cluster. It also provides a RESTful API for automation and advanced configuration scenarios.
- **Local Agent-based:** You can deploy a cluster locally with the Agent-based Installer for disconnected environments or restricted networks. It provides many of the benefits of the Assisted Installer, but you must download and configure the [Agent-based Installer](#) first. Configuration is done with a command-line interface. This approach is ideal for disconnected environments.
- **Automated:** You can deploy a cluster on installer-provisioned infrastructure. The installation program uses each cluster host's baseboard management controller (BMC) for provisioning. You can deploy clusters in connected or disconnected environments.
- **Full control:** You can deploy a cluster on infrastructure that you prepare and maintain, which provides maximum customizability. You can deploy clusters in connected or disconnected environments.

Each method deploys a cluster with the following characteristics:


- Highly available infrastructure with no single points of failure, which is available by default.
- Administrators can control what updates are applied and when.

1.1.1. About the installation program

You can use the installation program to deploy each type of cluster. The installation program generates the main assets, such as Ignition config files for the bootstrap, control plane, and compute machines. You can start an OpenShift Container Platform cluster with these three machine configurations, provided you correctly configured the infrastructure.

The OpenShift Container Platform installation program uses a set of targets and dependencies to manage cluster installations. The installation program has a set of targets that it must achieve, and each target has a set of dependencies. Because each target is only concerned with its own dependencies, the installation program can act to achieve multiple targets in parallel with the ultimate target being a running cluster. The installation program recognizes and uses existing components instead of running commands to create them again because the program meets the dependencies.

Figure 1.1. OpenShift Container Platform installation targets and dependencies



1.1.2. About Red Hat Enterprise Linux CoreOS (RHCOS)

Post-installation, each cluster machine uses Red Hat Enterprise Linux CoreOS (RHCOS) as the operating system. RHCOS is the immutable container host version of Red Hat Enterprise Linux (RHEL) and features a RHEL kernel with SELinux enabled by default. RHCOS includes the **kubelet**, which is the Kubernetes node agent, and the CRI-O container runtime, which is optimized for Kubernetes.

Every control plane machine in an OpenShift Container Platform 4.16 cluster must use RHCOS, which includes a critical first-boot provisioning tool called Ignition. This tool enables the cluster to configure the machines. Operating system updates are delivered as a bootable container image, using **OSTree** as a backend, that is deployed across the cluster by the Machine Config Operator. Actual operating system changes are made in-place on each machine as an atomic operation by using **rpm-ostree**. Together, these technologies enable OpenShift Container Platform to manage the operating system like it manages any other application on the cluster, by in-place upgrades that keep the entire platform up to date. These in-place updates can reduce the burden on operations teams.

If you use RHCOS as the operating system for all cluster machines, the cluster manages all aspects of its components and machines, including the operating system. Because of this, only the installation program and the Machine Config Operator can change machines. The installation program uses Ignition config files to set the exact state of each machine, and the Machine Config Operator completes more changes to the machines, such as the application of new certificates or keys, after installation.

1.1.3. Glossary of common terms for OpenShift Container Platform installing

The glossary defines common terms that relate to the installation content. Read the following list of terms to better understand the installation process.

Assisted Installer

An installer hosted at console.redhat.com that provides a web-based user interface or a RESTful API for creating a cluster configuration. The [Assisted Installer](#) generates a discovery image. Cluster machines boot with the discovery image, which installs RHCOS and an agent. Together, the Assisted Installer and agent provide preinstallation validation and installation for the cluster.

Agent-based Installer

An installer similar to the Assisted Installer, but you must download the [Agent-based Installer](#) first. The Agent-based Installer is ideal for disconnected environments.

Bootstrap node

A temporary machine that runs a minimal Kubernetes configuration required to deploy the OpenShift Container Platform control plane.

Control plane

A container orchestration layer that exposes the API and interfaces to define, deploy, and manage the lifecycle of containers. Also known as control plane machines.

Compute node

Nodes that are responsible for executing workloads for cluster users. Also known as worker nodes.

Disconnected installation

In some situations, parts of a data center might not have access to the internet, even through proxy servers. You can still install the OpenShift Container Platform in these environments, but you must download the required software and images and make them available to the disconnected environment.

The OpenShift Container Platform installation program

A program that provisions the infrastructure and deploys a cluster.

Installer-provisioned infrastructure

The installation program deploys and configures the infrastructure that the cluster runs on.

Ignition config files

A file that the Ignition tool uses to configure Red Hat Enterprise Linux CoreOS (RHCOS) during operating system initialization. The installation program generates different Ignition configuration files to initialize bootstrap, control plane, and worker nodes.

Kubernetes manifests

Specifications of a Kubernetes API object in a JSON or YAML format. A configuration file can include deployments, config maps, secrets, daemonsets, and so on.

Kubelet

A primary node agent that runs on each node in the cluster to ensure that containers are running in a pod.

Load balancers

A load balancer serves as the single point of contact for clients. Load balancers for the API distribute incoming traffic across control plane nodes.

Machine Config Operator

An Operator that manages and applies configurations and updates of the base operating system and container runtime, including everything between the kernel and kubelet, for the nodes in the cluster.

Operators

The preferred method of packaging, deploying, and managing a Kubernetes application in an OpenShift Container Platform cluster. An operator takes human operational knowledge and encodes it into software that is easily packaged and shared with customers.

User-provisioned infrastructure

You can install OpenShift Container Platform on infrastructure that you provide. You can use the installation program to generate the assets required to provision the cluster infrastructure, create the cluster infrastructure, and then deploy the cluster to the infrastructure that you provided.

1.1.4. Installation process

Except for the Assisted Installer, when you install an OpenShift Container Platform cluster, you must download the installation program from the appropriate [Cluster Type](#) page on the OpenShift Cluster Manager Hybrid Cloud Console. This console manages:

- REST API for accounts.
- Registry tokens, which are the pull secrets that you use to obtain the required components.
- Cluster registration, which associates the cluster identity to your Red Hat account to facilitate the gathering of usage metrics.

In OpenShift Container Platform 4.16, the installation program is a Go binary file that performs a series of file transformations on a set of assets. The way you interact with the installation program differs depending on your installation type. Consider the following installation use cases:

- To deploy a cluster with the Assisted Installer, you must configure the cluster settings by using the [Assisted Installer](#). There is no installation program to download and configure. After you finish setting the cluster configuration, you download a discovery ISO and then boot cluster machines with that image. You can install clusters with the Assisted Installer on Nutanix, vSphere, and bare metal with full integration, and other platforms without integration. If you install on bare metal, you must provide all of the cluster infrastructure and resources, including the networking, load balancing, storage, and individual cluster machines.
- To deploy clusters with the Agent-based Installer, you can download the [Agent-based Installer](#) first. You can then configure the cluster and generate a discovery image. You boot cluster machines with the discovery image, which installs an agent that communicates with the installation program and handles the provisioning for you instead of you interacting with the installation program or setting up a provisioner machine yourself. You must provide all of the cluster infrastructure and resources, including the networking, load balancing, storage, and individual cluster machines. This approach is ideal for disconnected environments.
- For clusters with installer-provisioned infrastructure, you delegate the infrastructure bootstrapping and provisioning to the installation program instead of doing it yourself. The installation program creates all of the networking, machines, and operating systems that are required to support the cluster, except if you install on bare metal. If you install on bare metal, you must provide all of the cluster infrastructure and resources, including the bootstrap machine, networking, load balancing, storage, and individual cluster machines.
- If you provision and manage the infrastructure for your cluster, you must provide all of the cluster infrastructure and resources, including the bootstrap machine, networking, load balancing, storage, and individual cluster machines.

For the installation program, the program uses three sets of files during installation: an installation configuration file that is named **install-config.yaml**, Kubernetes manifests, and Ignition config files for your machine types.



IMPORTANT

You can modify Kubernetes and the Ignition config files that control the underlying RHCOS operating system during installation. However, no validation is available to confirm the suitability of any modifications that you make to these objects. If you modify these objects, you might render your cluster non-functional. Because of this risk, modifying Kubernetes and Ignition config files is not supported unless you are following documented procedures or are instructed to do so by Red Hat support.

The installation configuration file is transformed into Kubernetes manifests, and then the manifests are wrapped into Ignition config files. The installation program uses these Ignition config files to create the cluster.

The installation configuration files are all pruned when you run the installation program, so be sure to back up all the configuration files that you want to use again.



IMPORTANT

You cannot modify the parameters that you set during installation, but you can modify many cluster attributes after installation.

The installation process with the Assisted Installer

Installation with the [Assisted Installer](#) involves creating a cluster configuration interactively by using the web-based user interface or the RESTful API. The Assisted Installer user interface prompts you for required values and provides reasonable default values for the remaining parameters, unless you change them in the user interface or with the API. The Assisted Installer generates a discovery image, which you download and use to boot the cluster machines. The image installs RHCOS and an agent, and the agent handles the provisioning for you. You can install OpenShift Container Platform with the Assisted Installer and full integration on Nutanix, vSphere, and bare metal. Additionally, you can install OpenShift Container Platform with the Assisted Installer on other platforms without integration.

OpenShift Container Platform manages all aspects of the cluster, including the operating system itself. Each machine boots with a configuration that references resources hosted in the cluster that it joins. This configuration allows the cluster to manage itself as updates are applied.

If possible, use the Assisted Installer feature to avoid having to download and configure the Agent-based Installer.

The installation process with Agent-based infrastructure

Agent-based installation is similar to using the Assisted Installer, except that you must initially download and install the [Agent-based Installer](#). An Agent-based installation is useful when you want the convenience of the Assisted Installer, but you need to install a cluster in a disconnected environment.

If possible, use the Agent-based installation feature to avoid having to create a provisioner machine with a bootstrap VM, and then provision and maintain the cluster infrastructure.

The installation process with installer-provisioned infrastructure

The default installation type uses installer-provisioned infrastructure. By default, the installation program acts as an installation wizard, prompting you for values that it cannot determine on its own and providing reasonable default values for the remaining parameters. You can also customize the installation process to support advanced infrastructure scenarios. The installation program provisions the underlying infrastructure for the cluster.

You can install either a standard cluster or a customized cluster. With a standard cluster, you provide minimum details that are required to install the cluster. With a customized cluster, you can specify more details about the platform, such as the number of machines that the control plane uses, the type of

virtual machine that the cluster deploys, or the CIDR range for the Kubernetes service network.

If possible, use this feature to avoid having to provision and maintain the cluster infrastructure. In all other environments, you use the installation program to generate the assets that you require to provision your cluster infrastructure.

With installer-provisioned infrastructure clusters, OpenShift Container Platform manages all aspects of the cluster, including the operating system itself. Each machine boots with a configuration that references resources hosted in the cluster that it joins. This configuration allows the cluster to manage itself as updates are applied.

The installation process with user-provisioned infrastructure

You can also install OpenShift Container Platform on infrastructure that you provide. You use the installation program to generate the assets that you require to provision the cluster infrastructure, create the cluster infrastructure, and then deploy the cluster to the infrastructure that you provided.

If you do not use infrastructure that the installation program provisioned, you must manage and maintain the cluster resources yourself. The following list details some of these self-managed resources:


- The underlying infrastructure for the control plane and compute machines that make up the cluster
- Load balancers
- Cluster networking, including the DNS records and required subnets
- Storage for the cluster infrastructure and applications

If your cluster uses user-provisioned infrastructure, you have the option of adding RHEL compute machines to your cluster.

Installation process details

When a cluster is provisioned, each machine in the cluster requires information about the cluster. OpenShift Container Platform uses a temporary bootstrap machine during initial configuration to provide the required information to the permanent control plane. The temporary bootstrap machine boots by using an Ignition config file that describes how to create the cluster. The bootstrap machine creates the control plane machines that make up the control plane. The control plane machines then create the compute machines, which are also known as worker machines. The following figure illustrates this process:

Figure 1.2. Creating the bootstrap, control plane, and compute machines



After the cluster machines initialize, the bootstrap machine is destroyed. All clusters use the bootstrap process to initialize the cluster, but if you provision the infrastructure for your cluster, you must complete many of the steps manually.

IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- Consider using Ignition config files within 12 hours after they are generated, because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

Bootstrapping a cluster involves the following steps:

1. The bootstrap machine boots and starts hosting the remote resources required for the control plane machines to boot. If you provision the infrastructure, this step requires manual intervention.
2. The bootstrap machine starts a single-node etcd cluster and a temporary Kubernetes control plane.
3. The control plane machines fetch the remote resources from the bootstrap machine and finish booting. If you provision the infrastructure, this step requires manual intervention.

4. The temporary control plane schedules the production control plane to the production control plane machines.
5. The Cluster Version Operator (CVO) comes online and installs the etcd Operator. The etcd Operator scales up etcd on all control plane nodes.
6. The temporary control plane shuts down and passes control to the production control plane.
7. The bootstrap machine injects OpenShift Container Platform components into the production control plane.
8. The installation program shuts down the bootstrap machine. If you provision the infrastructure, this step requires manual intervention.
9. The control plane sets up the compute nodes.
10. The control plane installs additional services in the form of a set of Operators.

The result of this bootstrapping process is a running OpenShift Container Platform cluster. The cluster then downloads and configures remaining components needed for the day-to-day operations, including the creation of compute machines in supported environments.

Additional resources

- [Red Hat OpenShift Network Calculator](#)

1.1.5. Verifying node state after installation

The OpenShift Container Platform installation completes when the following installation health checks are successful:

- The provisioner can access the OpenShift Container Platform web console.
- All control plane nodes are ready.
- All cluster Operators are available.



NOTE

After the installation completes, the specific cluster Operators responsible for the worker nodes continuously attempt to provision all worker nodes. Some time is required before all worker nodes report as **READY**. For installations on bare metal, wait a minimum of 60 minutes before troubleshooting a worker node. For installations on all other platforms, wait a minimum of 40 minutes before troubleshooting a worker node. A **DEGRADED** state for the cluster Operators responsible for the worker nodes depends on the Operators' own resources and not on the state of the nodes.

After your installation completes, you can continue to monitor the condition of the nodes in your cluster.

Prerequisites

- The installation program resolves successfully in the terminal.

Procedure

1. Show the status of all worker nodes:

```
$ oc get nodes
```

Example output

```
NAME                                STATUS ROLES  AGE  VERSION
example-compute1.example.com       Ready  worker  13m  v1.21.6+bb8d50a
example-compute2.example.com       Ready  worker  13m  v1.21.6+bb8d50a
example-compute4.example.com       Ready  worker  14m  v1.21.6+bb8d50a
example-control1.example.com       Ready  master  52m  v1.21.6+bb8d50a
example-control2.example.com       Ready  master  55m  v1.21.6+bb8d50a
example-control3.example.com       Ready  master  55m  v1.21.6+bb8d50a
```

2. Show the phase of all worker machine nodes:

```
$ oc get machines -A
```

Example output

```
NAMESPACE      NAME                                PHASE  TYPE  REGION  ZONE  AGE
openshift-machine-api  example-zbbt6-master-0             Running  Running  Running  95m
openshift-machine-api  example-zbbt6-master-1             Running  Running  Running  95m
openshift-machine-api  example-zbbt6-master-2             Running  Running  Running  95m
openshift-machine-api  example-zbbt6-worker-0-25bhp       Running  Running  Running  49m
openshift-machine-api  example-zbbt6-worker-0-8b4c2       Running  Running  Running  49m
openshift-machine-api  example-zbbt6-worker-0-jkbqt       Running  Running  Running  49m
openshift-machine-api  example-zbbt6-worker-0-qrl5b       Running  Running  Running  49m
```

Additional resources

- [Getting the BareMetalHost resource](#)
- [Following the installation](#)
- [Validating an installation](#)
- [Agent-based Installer](#)
- [Assisted Installer for OpenShift Container Platform](#)

Installation scope

The scope of the OpenShift Container Platform installation program is intentionally narrow. It is designed for simplicity and ensured success. You can complete many more configuration tasks after installation completes.

Additional resources

- See [Available cluster customizations](#) for details about OpenShift Container Platform configuration resources.

1.1.6. OpenShift Local overview

OpenShift Local supports rapid application development to get started building OpenShift Container Platform clusters. OpenShift Local is designed to run on a local computer to simplify setup and testing, and to emulate the cloud development environment locally with all of the tools needed to develop container-based applications.

Regardless of the programming language you use, OpenShift Local hosts your application and brings a minimal, preconfigured Red Hat OpenShift Container Platform cluster to your local PC without the need for a server-based infrastructure.

On a hosted environment, OpenShift Local can create microservices, convert them into images, and run them in Kubernetes-hosted containers directly on your laptop or desktop running Linux, macOS, or Windows 10 or later.


For more information about OpenShift Local, see [Red Hat OpenShift Local Overview](#).

1.2. SUPPORTED PLATFORMS FOR OPENSIFT CONTAINER PLATFORM CLUSTERS

In OpenShift Container Platform 4.16, you can install a cluster that uses installer-provisioned infrastructure on the following platforms:

- Alibaba Cloud
- Amazon Web Services (AWS)
- Bare metal
- Google Cloud Platform (GCP)
- IBM Cloud®
- Microsoft Azure
- Microsoft Azure Stack Hub
- Nutanix
- Red Hat OpenStack Platform (RHOSP)
 - The latest OpenShift Container Platform release supports both the latest RHOSP long-life release and intermediate release. For complete RHOSP release compatibility, see the [OpenShift Container Platform on RHOSP support matrix](#).
- VMware vSphere

For these clusters, all machines, including the computer that you run the installation process on, must have direct internet access to pull images for platform containers and provide telemetry data to Red Hat.



IMPORTANT

After installation, the following changes are not supported:

- Mixing cloud provider platforms.
- Mixing cloud provider components. For example, using a persistent storage framework from a another platform on the platform where you installed the cluster.

In OpenShift Container Platform 4.16, you can install a cluster that uses user-provisioned infrastructure on the following platforms:

- AWS
- Azure
- Azure Stack Hub
- Bare metal
- GCP
- IBM Power®
- IBM Z® or IBM® LinuxONE
- RHOSP
 - The latest OpenShift Container Platform release supports both the latest RHOSP long-life release and intermediate release. For complete RHOSP release compatibility, see the [OpenShift Container Platform on RHOSP support matrix](#).
- VMware Cloud on AWS
- VMware vSphere

Depending on the supported cases for the platform, you can perform installations on user-provisioned infrastructure, so that you can run machines with full internet access, place your cluster behind a proxy, or perform a disconnected installation.

In a disconnected installation, you can download the images that are required to install a cluster, place them in a mirror registry, and use that data to install your cluster. While you require internet access to pull images for platform containers, with a disconnected installation on vSphere or bare metal infrastructure, your cluster machines do not require direct internet access.

The [OpenShift Container Platform 4.x Tested Integrations](#) page contains details about integration testing for different platforms.

Additional resources

- See [Supported installation methods for different platforms](#) for more information about the types of installations that are available for each supported platform.
- See [Selecting a cluster installation method and preparing it for users](#) for information about choosing an installation method and preparing the required resources.

- [Red Hat OpenShift Network Calculator](#) can help you design your cluster network during both the deployment and expansion phases. It addresses common questions related to the cluster network and provides output in a convenient JSON format.

CHAPTER 2. SELECTING A CLUSTER INSTALLATION METHOD AND PREPARING IT FOR USERS

Before you install OpenShift Container Platform, decide what kind of installation process to follow and verify that you have all of the required resources to prepare the cluster for users.

2.1. SELECTING A CLUSTER INSTALLATION TYPE

Before you install an OpenShift Container Platform cluster, you need to select the best installation instructions to follow. Think about your answers to the following questions to select the best option.

2.1.1. Do you want to install and manage an OpenShift Container Platform cluster yourself?

If you want to install and manage OpenShift Container Platform yourself, you can install it on the following platforms:

- Amazon Web Services (AWS) on 64-bit x86 instances
- Amazon Web Services (AWS) on 64-bit ARM instances
- Microsoft Azure on 64-bit x86 instances
- Microsoft Azure on 64-bit ARM instances
- Microsoft Azure Stack Hub
- Google Cloud Platform (GCP) on 64-bit x86 instances
- Google Cloud Platform (GCP) on 64-bit ARM instances
- Red Hat OpenStack Platform (RHOSP)
- IBM Cloud®
- IBM Z® or IBM® LinuxONE
- IBM Z® or IBM® LinuxONE for Red Hat Enterprise Linux (RHEL) KVM
- IBM Power®
- IBM Power® Virtual Server
- Nutanix
- VMware vSphere
- Bare metal or other platform agnostic infrastructure

You can deploy an OpenShift Container Platform 4 cluster to both on-premise hardware and to cloud hosting services, but all of the machines in a cluster must be in the same data center or cloud hosting service.

If you want to use OpenShift Container Platform but do not want to manage the cluster yourself, you have several managed service options. If you want a cluster that is fully managed by Red Hat, you can use

[OpenShift Dedicated](#) or [OpenShift Online](#). You can also use OpenShift as a managed service on Azure, AWS, IBM Cloud®, or Google Cloud. For more information about managed services, see the [OpenShift Products](#) page. If you install an OpenShift Container Platform cluster with a cloud virtual machine as a virtual bare metal, the corresponding cloud-based storage is not supported.

2.1.2. Have you used OpenShift Container Platform 3 and want to use OpenShift Container Platform 4?

If you used OpenShift Container Platform 3 and want to try OpenShift Container Platform 4, you need to understand how different OpenShift Container Platform 4 is. OpenShift Container Platform 4 weaves the Operators that package, deploy, and manage Kubernetes applications and the operating system that the platform runs on, Red Hat Enterprise Linux CoreOS (RHCOS), together seamlessly. Instead of deploying machines and configuring their operating systems so that you can install OpenShift Container Platform on them, the RHCOS operating system is an integral part of the OpenShift Container Platform cluster. Deploying the operating system for the cluster machines is part of the installation process for OpenShift Container Platform. See [Differences between OpenShift Container Platform 3 and 4](#).

Because you need to provision machines as part of the OpenShift Container Platform cluster installation process, you cannot upgrade an OpenShift Container Platform 3 cluster to OpenShift Container Platform 4. Instead, you must create a new OpenShift Container Platform 4 cluster and migrate your OpenShift Container Platform 3 workloads to them. For more information about migrating, see [Migrating from OpenShift Container Platform 3 to 4 overview](#). Because you must migrate to OpenShift Container Platform 4, you can use any type of production cluster installation process to create your new cluster.

2.1.3. Do you want to use existing components in your cluster?

Because the operating system is integral to OpenShift Container Platform, it is easier to let the installation program for OpenShift Container Platform stand up all of the infrastructure. These are called *installer provisioned infrastructure* installations. In this type of installation, you can provide some existing infrastructure to the cluster, but the installation program deploys all of the machines that your cluster initially needs.

You can deploy an installer-provisioned infrastructure cluster without specifying any customizations to the cluster or its underlying machines to [AWS](#), [Azure](#), [Azure Stack Hub](#), [GCP](#), [Nutanix](#).

If you need to perform basic configuration for your installer-provisioned infrastructure cluster, such as the instance type for the cluster machines, you can customize an installation for [AWS](#), [Azure](#), [GCP](#), [Nutanix](#).

For installer-provisioned infrastructure installations, you can use an existing [VPC in AWS](#), [vNet in Azure](#), or [VPC in GCP](#). You can also reuse part of your networking infrastructure so that your cluster in [AWS](#), [Azure](#), [GCP](#) can coexist with existing IP address allocations in your environment and integrate with existing MTU and VXLAN configurations. If you have existing accounts and credentials on these clouds, you can re-use them, but you might need to modify the accounts to have the required permissions to install OpenShift Container Platform clusters on them.

You can use the installer-provisioned infrastructure method to create appropriate machine instances on your hardware for [vSphere](#), and [bare metal](#). Additionally, for [vSphere](#), you can also customize additional network parameters during installation.

If you want to reuse extensive cloud infrastructure, you can complete a *user-provisioned infrastructure* installation. With these installations, you manually deploy the machines that your cluster requires during the installation process. If you perform a user-provisioned infrastructure installation on [AWS](#), [Azure](#),

[Azure Stack Hub](#), you can use the provided templates to help you stand up all of the required components. You can also reuse a shared [VPC on GCP](#). Otherwise, you can use the [provider-agnostic](#) installation method to deploy a cluster into other clouds.

You can also complete a user-provisioned infrastructure installation on your existing hardware. If you use [RHOSP](#), [IBM Z® or IBM® LinuxONE](#), [IBM Z® and IBM® LinuxONE with RHEL KVM](#), [IBM Power](#), or [vSphere](#), use the specific installation instructions to deploy your cluster. If you use other supported hardware, follow the [bare metal installation](#) procedure. For some of these platforms, such as [vSphere](#), and [bare metal](#), you can also customize additional network parameters during installation.

2.1.4. Do you need extra security for your cluster?

If you use a user-provisioned installation method, you can configure a proxy for your cluster. The instructions are included in each installation procedure.

If you want to prevent your cluster on a public cloud from exposing endpoints externally, you can deploy a private cluster with installer-provisioned infrastructure on [AWS](#), [Azure](#), or [GCP](#).

If you need to install your cluster that has limited access to the internet, such as a disconnected or restricted network cluster, you can [mirror the installation packages](#) and install the cluster from them. Follow detailed instructions for user-provisioned infrastructure installations into restricted networks for [AWS](#), [GCP](#), [IBM Z® or IBM® LinuxONE](#), [IBM Z® or IBM® LinuxONE with RHEL KVM](#), [IBM Power®](#), [vSphere](#), or [bare metal](#). You can also install a cluster into a restricted network using installer-provisioned infrastructure by following detailed instructions for [AWS](#), [GCP](#), [IBM Cloud®](#), [Nutanix](#), [RHOSP](#), and [vSphere](#).

If you need to deploy your cluster to an [AWS GovCloud region](#), [AWS China region](#), or [Azure government region](#), you can configure those custom regions during an installer-provisioned infrastructure installation.

You can also configure the cluster machines to use the RHEL cryptographic libraries that have been submitted to NIST for [FIPS 140-2/140-3 Validation](#) during installation.



IMPORTANT

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

2.2. PREPARING YOUR CLUSTER FOR USERS AFTER INSTALLATION

Some configuration is not required to install the cluster but recommended before your users access the cluster. You can customize the cluster itself by [customizing](#) the Operators that make up your cluster and integrate your cluster with other required systems, such as an identity provider.

For a production cluster, you must configure the following integrations:

- [Persistent storage](#)
- [An identity provider](#)
- [Monitoring core OpenShift Container Platform components](#)

2.3. PREPARING YOUR CLUSTER FOR WORKLOADS

Depending on your workload needs, you might need to take extra steps before you begin deploying applications. For example, after you prepare infrastructure to support your application [build strategy](#), you might need to make provisions for [low-latency](#) workloads or to [protect sensitive workloads](#). You can also configure [monitoring](#) for application workloads.

2.4. SUPPORTED INSTALLATION METHODS FOR DIFFERENT PLATFORMS

You can perform different types of installations on different platforms.



NOTE

Not all installation options are supported for all platforms, as shown in the following tables. A checkmark indicates that the option is supported and links to the relevant section.

Table 2.1. Installer-provisioned infrastructure options

	AWS	Azure	Azure	Azure	Google Cloud	Google Cloud	Nutanix	RHEL	Bare Metal	Bare Metal	vSphere	IBM Cloud	IBM Z	IBM Power	IBM Power
Default	✓	✓	✓	✓	✓	✓	✓		✓	✓	✓	✓			
Custom	✓	✓	✓	✓	✓	✓	✓	✓			✓	✓			✓
Network customization	✓	✓	✓	✓	✓	✓					✓	✓			

	AWS (64-bit x86_64)	AWS (64-bit ARM)	Azure (64-bit x86_64)	Azure (64-bit ARM)	Azure (Stable Hub)	Google Cloud (64-bit x86_64)	Google Cloud (64-bit ARM)	Nutanix	RHEL OS P	Baremetal (64-bit x86_64)	Baremetal (64-bit ARM)	vSphere	IBM Cloud®	IBM Z®	IBM Power®	IBM Power® Virtual Server
Restricted network	✓	✓	✓	✓		✓	✓	✓	✓	✓	✓	✓	✓			✓
Private clusters	✓	✓	✓	✓		✓	✓						✓			✓
Existing virtual private networks	✓	✓	✓	✓		✓	✓						✓			✓

	AWS (64-bit x86_64)	Azure (64-bit x86_64)	Azure (64-bit ARM)	Azure (64-bit ARM)	Google Cloud (64-bit x86_64)	Google Cloud (64-bit ARM)	Nutanix	RHEL P	Baremetal (64-bit x86_64)	Baremetal (64-bit ARM)	vSphere	IBM Cloud®	IBM Z®	IBM Power®	IBM Power® Virtual Server		
Governments	✓		✓														
Secret regions	✓																
China regions	✓																

Table 2.2. User-provisioned infrastructure options

	AWS (64-bit x86_64)	Azure (64-bit ARM)	Azure (64-bit ARM)	Azure (64-bit ARM)	Google Cloud (64-bit x86_64)	Google Cloud (64-bit ARM)	Nutanix	RHO SP	Baremetal (64-bit x86_64)	Baremetal (64-bit ARM)	vSphere	IBM Cloud®	IBM Z®	IBM Z® with RHEL KVM	IBM Power®	Platform agnostic
Custom	✓	✓	✓	✓	✓	✓		✓	✓	✓	✓		✓	✓	✓	✓
Network customization									✓	✓	✓					
Restricted network	✓	✓				✓	✓		✓	✓	✓		✓	✓	✓	

	AWS (64-bit x86_64)	Azure (64-bit x86_64)	Azure (64-bit ARM)	Azure (64-bit ARM)	Google Cloud (64-bit x86_64)	Google Cloud (64-bit ARM)	Nutanix	RHO SP	Baremetal (64-bit x86_64)	Baremetal (64-bit ARM)	vSphere	IBM Cloud®	IBM Z®	IBM Z® with RHEL KVM	IBM Power®	Platform agnostic
Shared VPC hosted outside of cluster project					✓	✓										

CHAPTER 3. CLUSTER CAPABILITIES

Cluster administrators can use cluster capabilities to enable or disable optional components prior to installation. Cluster administrators can enable cluster capabilities at anytime after installation.




NOTE

Cluster administrators cannot disable a cluster capability after it is enabled.

3.1. SELECTING CLUSTER CAPABILITIES

You can select cluster capabilities by following one of the installation methods that include customizing your cluster, such as "Installing a cluster on AWS with customizations" or "Installing a cluster on GCP with customizations".

During a customized installation, you create an **install-config.yaml** file that contains the configuration parameters for your cluster.



NOTE


If you customize your cluster by enabling or disabling specific cluster capabilities, you are responsible for manually maintaining your **install-config.yaml** file. New OpenShift Container Platform updates might declare new capability handles for existing components, or introduce new components altogether. Users who customize their **install-config.yaml** file should consider periodically updating their **install-config.yaml** file as OpenShift Container Platform is updated.

You can use the following configuration parameters to select cluster capabilities:

capabilities:

```
baselineCapabilitySet: v4.11 1
additionalEnabledCapabilities: 2
- CSISnapshot
- Console
- Storage
```


- 1** Defines a baseline set of capabilities to install. Valid values are **None**, **vCurrent** and **v4.x**. If you select **None**, all optional capabilities will be disabled. The default value is **vCurrent**, which enables all optional capabilities.



NOTE

v4.x refers to any value up to and including the current cluster version. For example, valid values for a OpenShift Container Platform 4.12 cluster are **v4.11** and **v4.12**.

- 2** Defines a list of capabilities to explicitly enable. These will be enabled in addition to the capabilities specified in **baselineCapabilitySet**.

**NOTE**

In this example, the default capability is set to **v4.11**. The **additionalEnabledCapabilities** field enables additional capabilities over the default **v4.11** capability set.

The following table describes the **baselineCapabilitySet** values.

Table 3.1. Cluster capabilities **baselineCapabilitySet values description**

Value	Description
vCurrent	Specify this option when you want to automatically add new, default capabilities that are introduced in new releases.
v4.11	Specify this option when you want to enable the default capabilities for OpenShift Container Platform 4.11. By specifying v4.11 , capabilities that are introduced in newer versions of OpenShift Container Platform are not enabled. The default capabilities in OpenShift Container Platform 4.11 are baremetal, MachineAPI, marketplace, and openshift-samples .
v4.12	Specify this option when you want to enable the default capabilities for OpenShift Container Platform 4.12. By specifying v4.12 , capabilities that are introduced in newer versions of OpenShift Container Platform are not enabled. The default capabilities in OpenShift Container Platform 4.12 are baremetal, MachineAPI, marketplace, openshift-samples, Console, Insights, Storage, and CSISnapshot .
v4.13	Specify this option when you want to enable the default capabilities for OpenShift Container Platform 4.13. By specifying v4.13 , capabilities that are introduced in newer versions of OpenShift Container Platform are not enabled. The default capabilities in OpenShift Container Platform 4.13 are baremetal, MachineAPI, marketplace, openshift-samples, Console, Insights, Storage, CSISnapshot, and NodeTuning .
v4.14	Specify this option when you want to enable the default capabilities for OpenShift Container Platform 4.14. By specifying v4.14 , capabilities that are introduced in newer versions of OpenShift Container Platform are not enabled. The default capabilities in OpenShift Container Platform 4.14 are baremetal, MachineAPI, marketplace, openshift-samples, Console, Insights, Storage, CSISnapshot, NodeTuning, ImageRegistry, Build, and DeploymentConfig .

Value	Description
v4.15	Specify this option when you want to enable the default capabilities for OpenShift Container Platform 4.15. By specifying v4.15 , capabilities that are introduced in newer versions of OpenShift Container Platform are not enabled. The default capabilities in OpenShift Container Platform 4.15 are baremetal, MachineAPI, marketplace, OperatorLifecycleManager, openshift-samples, Console, Insights, Storage, CSISnapshot, NodeTuning, ImageRegistry, Build, CloudCredential , and DeploymentConfig .
v4.16	Specify this option when you want to enable the default capabilities for OpenShift Container Platform 4.16. By specifying v4.16 , capabilities that are introduced in newer versions of OpenShift Container Platform are not enabled. The default capabilities in OpenShift Container Platform 4.16 are baremetal, MachineAPI, marketplace, OperatorLifecycleManager, openshift-samples, Console, Insights, Storage, CSISnapshot, NodeTuning, ImageRegistry, Build, CloudCredential, DeploymentConfig , and CloudControllerManager .
None	Specify when the other sets are too large, and you do not need any capabilities or want to fine-tune via additionalEnabledCapabilities .

Additional resources

- [Installing a cluster on AWS with customizations](#)
- [Installing a cluster on GCP with customizations](#)

3.2. OPTIONAL CLUSTER CAPABILITIES IN OPENSIFT CONTAINER PLATFORM 4.16

Currently, cluster Operators provide the features for these optional capabilities. The following summarizes the features provided by each capability and what functionality you lose if it is disabled.

Additional resources

- [Cluster Operators reference](#)

3.2.1. Bare-metal capability

Purpose


The Cluster Baremetal Operator provides the features for the **baremetal** capability.

The Cluster Baremetal Operator (CBO) deploys all the components necessary to take a bare-metal server to a fully functioning worker node ready to run OpenShift Container Platform compute nodes. The CBO ensures that the metal3 deployment, which consists of the Bare Metal Operator (BMO) and

Ironic containers, runs on one of the control plane nodes within the OpenShift Container Platform cluster. The CBO also listens for OpenShift Container Platform updates to resources that it watches and takes appropriate action.

The bare-metal capability is required for deployments using installer-provisioned infrastructure. Disabling the bare-metal capability can result in unexpected problems with these deployments.

It is recommended that cluster administrators only disable the bare-metal capability during installations with user-provisioned infrastructure that do not have any **BareMetalHost** resources in the cluster.



IMPORTANT

If the bare-metal capability is disabled, the cluster cannot provision or manage bare-metal nodes. Only disable the capability if there are no **BareMetalHost** resources in your deployment. The **baremetal** capability depends on the **MachineAPI** capability. If you enable the **baremetal** capability, you must also enable **MachineAPI**.

Additional resources

- [Deploying installer-provisioned clusters on bare metal](#)
- [Preparing for bare metal cluster installation](#)
- [Bare metal configuration](#)

3.2.2. Build capability

Purpose

The **Build** capability enables the **Build** API. The **Build** API manages the lifecycle of **Build** and **BuildConfig** objects.



IMPORTANT

If you disable the **Build** capability, the following resources will not be available in the cluster:

- **Build** and **BuildConfig** resources
- The **builder** service account

Disable the **Build** capability only if you do not require **Build** and **BuildConfig** resources or the **builder** service account in the cluster.

3.2.3. Cloud controller manager capability

Purpose

The Cloud Controller Manager Operator provides features for the **CloudControllerManager** capability.



NOTE

Currently, disabling the **CloudControllerManager** capability is not supported on all platforms.

You can determine if your cluster supports disabling the **CloudControllerManager** capability by checking values in the installation configuration (**install-config.yaml**) file for your cluster.

In the **install-config.yaml** file, locate the **platform** parameter.

- If the value of the **platform** parameter is **Baremetal** or **None**, you can disable the **CloudControllerManager** capability on your cluster.
- If the value of the **platform** parameter is **External**, locate the **platform.external.cloudControllerManager** parameter. If the value of the **platform.external.cloudControllerManager** parameter is **None**, you can disable the **CloudControllerManager** capability on your cluster.



IMPORTANT

If these parameters contain any other values than those listed, you cannot disable the **CloudControllerManager** capability on your cluster.



NOTE

The status of this Operator is General Availability for Amazon Web Services (AWS), Google Cloud Platform (GCP), IBM Cloud®, global Microsoft Azure, Microsoft Azure Stack Hub, Nutanix, Red Hat OpenStack Platform (RHOSP), and VMware vSphere.

The Operator is available as a [Technology Preview](#) for Alibaba Cloud and IBM Power® Virtual Server.

The Cloud Controller Manager Operator manages and updates the cloud controller managers deployed on top of OpenShift Container Platform. The Operator is based on the Kubebuilder framework and **controller-runtime** libraries. It is installed via the Cluster Version Operator (CVO).

It contains the following components:

- Operator
- Cloud configuration observer

By default, the Operator exposes Prometheus metrics through the **metrics** service.

3.2.4. Cloud credential capability

Purpose

The Cloud Credential Operator provides features for the **CloudCredential** capability.



NOTE

Currently, disabling the **CloudCredential** capability is only supported for bare-metal clusters.

The Cloud Credential Operator (CCO) manages cloud provider credentials as Kubernetes custom resource definitions (CRDs). The CCO syncs on **CredentialsRequest** custom resources (CRs) to allow OpenShift Container Platform components to request cloud provider credentials with the specific permissions that are required for the cluster to run.

By setting different values for the **credentialsMode** parameter in the **install-config.yaml** file, the CCO can be configured to operate in several different modes. If no mode is specified, or the **credentialsMode** parameter is set to an empty string (""), the CCO operates in its default mode.

Additional resources

- [About the Cloud Credential Operator](#)

3.2.5. Cluster storage capability

Purpose

The Cluster Storage Operator provides the features for the **Storage** capability.

The Cluster Storage Operator sets OpenShift Container Platform cluster-wide storage defaults. It ensures a default **storageclass** exists for OpenShift Container Platform clusters. It also installs Container Storage Interface (CSI) drivers which enable your cluster to use various storage backends.



IMPORTANT

If the cluster storage capability is disabled, the cluster will not have a default **storageclass** or any CSI drivers. Users with administrator privileges can create a default **storageclass** and manually install CSI drivers if the cluster storage capability is disabled.

Notes

- The storage class that the Operator creates can be made non-default by editing its annotation, but this storage class cannot be deleted as long as the Operator runs.

3.2.6. Console capability

Purpose

The Console Operator provides the features for the **Console** capability.

The Console Operator installs and maintains the OpenShift Container Platform web console on a cluster. The Console Operator is installed by default and automatically maintains a console.

Additional resources

- [Web console overview](#)

3.2.7. CSI snapshot controller capability

Purpose

The Cluster CSI Snapshot Controller Operator provides the features for the **CSISnapshot** capability.

The Cluster CSI Snapshot Controller Operator installs and maintains the CSI Snapshot Controller. The CSI Snapshot Controller is responsible for watching the **VolumeSnapshot** CRD objects and manages the creation and deletion lifecycle of volume snapshots.

Additional resources

- [CSI volume snapshots](#)

3.2.8. DeploymentConfig capability

Purpose

The **DeploymentConfig** capability enables and manages the **DeploymentConfig** API.

**IMPORTANT**

If you disable the **DeploymentConfig** capability, the following resources will not be available in the cluster:

- **DeploymentConfig** resources
- The **deployer** service account

Disable the **DeploymentConfig** capability only if you do not require **DeploymentConfig** resources and the **deployer** service account in the cluster.

3.2.9. Ingress Capability**Purpose**

The Ingress Operator provides the features for the **Ingress** capability.

The Ingress Operator configures and manages the OpenShift Container Platform router.

Project

[openshift-ingress-operator](#)

CRDs

- **clusteringresses.ingress.openshift.io**
 - Scope: Namespaced
 - CR: **clusteringresses**
 - Validation: No

Configuration objects

- Cluster config
 - Type Name: **clusteringresses.ingress.openshift.io**
 - Instance Name: **default**
 - View Command:

```
$ oc get clusteringresses.ingress.openshift.io -n openshift-ingress-operator default -o yaml
```

Notes

The Ingress Operator sets up the router in the **openshift-ingress** project and creates the deployment for the router:

```
$ oc get deployment -n openshift-ingress
```

The Ingress Operator uses the **clusterNetwork[].cidr** from the **network/cluster** status to determine what mode (IPv4, IPv6, or dual stack) the managed Ingress Controller (router) should operate in. For

example, if **clusterNetwork** contains only a v6 **cidr**, then the Ingress Controller operates in IPv6-only mode.

In the following example, Ingress Controllers managed by the Ingress Operator will run in IPv4-only mode because only one cluster network exists and the network is an IPv4 **cidr**:

```
$ oc get network/cluster -o jsonpath='{.status.clusterNetwork[*]}'
```

Example output

```
map[cidr:10.128.0.0/14 hostPrefix:23]
```

3.2.10. Insights capability

Purpose

The Insights Operator provides the features for the **Insights** capability.

The Insights Operator gathers OpenShift Container Platform configuration data and sends it to Red Hat. The data is used to produce proactive insights recommendations about potential issues that a cluster might be exposed to. These insights are communicated to cluster administrators through Insights Advisor on console.redhat.com.

Notes

Insights Operator complements OpenShift Container Platform Telemetry.

Additional resources

- [Using Insights Operator](#)

3.2.11. Machine API capability

Purpose

The **machine-api-operator**, **cluster-autoscaler-operator**, and **cluster-control-plane-machine-set-operator** Operators provide the features for the **MachineAPI** capability. You can disable this capability only if you install a cluster with user-provisioned infrastructure.

The Machine API capability is responsible for all machine configuration and management in the cluster. If you disable the Machine API capability during installation, you need to manage all machine-related tasks manually.

Additional resources

- [Overview of machine management](#)
- [Machine API Operator](#)
- [Cluster Autoscaler Operator](#)
- [Control Plane Machine Set Operator](#)

3.2.12. Marketplace capability

Purpose

The Marketplace Operator provides the features for the **marketplace** capability.

The Marketplace Operator simplifies the process for bringing off-cluster Operators to your cluster by using a set of default Operator Lifecycle Manager (OLM) catalogs on the cluster. When the Marketplace Operator is installed, it creates the **openshift-marketplace** namespace. OLM ensures catalog sources installed in the **openshift-marketplace** namespace are available for all namespaces on the cluster.

If you disable the **marketplace** capability, the Marketplace Operator does not create the **openshift-marketplace** namespace. Catalog sources can still be configured and managed on the cluster manually, but OLM depends on the **openshift-marketplace** namespace in order to make catalogs available to all namespaces on the cluster. Users with elevated permissions to create namespaces prefixed with **openshift-**, such as system or cluster administrators, can manually create the **openshift-marketplace** namespace.

If you enable the **marketplace** capability, you can enable and disable individual catalogs by configuring the Marketplace Operator.

Additional resources

- [Red Hat-provided Operator catalogs](#)

3.2.13. Operator Lifecycle Manager capability

Purpose

Operator Lifecycle Manager (OLM) helps users install, update, and manage the lifecycle of Kubernetes native applications (Operators) and their associated services running across their OpenShift Container Platform clusters. It is part of the [Operator Framework](#), an open source toolkit designed to manage Operators in an effective, automated, and scalable way.

If an Operator requires any of the following APIs, then you must enable the **OperatorLifecycleManager** capability:

- **ClusterServiceVersion**
- **CatalogSource**
- **Subscription**
- **InstallPlan**
- **OperatorGroup**



IMPORTANT

The **marketplace** capability depends on the **OperatorLifecycleManager** capability. You cannot disable the **OperatorLifecycleManager** capability and enable the **marketplace** capability.

Additional resources

- [Operator Lifecycle Manager concepts and resources](#)

3.2.14. Node Tuning capability

Purpose

The Node Tuning Operator provides features for the **NodeTuning** capability.

The Node Tuning Operator helps you manage node-level tuning by orchestrating the TuneD daemon and achieves low latency performance by using the Performance Profile controller. The majority of high-performance applications require some level of kernel tuning. The Node Tuning Operator provides a unified management interface to users of node-level sysctls and more flexibility to add custom tuning specified by user needs.

If you disable the NodeTuning capability, some default tuning settings will not be applied to the control-plane nodes. This might limit the scalability and performance of large clusters with over 900 nodes or 900 routes.

Additional resources

- [Using the Node Tuning Operator](#)

3.2.15. OpenShift samples capability

Purpose

The Cluster Samples Operator provides the features for the **openshift-samples** capability.

The Cluster Samples Operator manages the sample image streams and templates stored in the **openshift** namespace.

On initial start up, the Operator creates the default samples configuration resource to initiate the creation of the image streams and templates. The configuration object is a cluster scoped object with the key **cluster** and type **configs.samples**.

The image streams are the Red Hat Enterprise Linux CoreOS (RHCOS)-based OpenShift Container Platform image streams pointing to images on **registry.redhat.io**. Similarly, the templates are those categorized as OpenShift Container Platform templates.

If you disable the samples capability, users cannot access the image streams, samples, and templates it provides. Depending on your deployment, you might want to disable this component if you do not need it.

Additional resources

- [Configuring the Cluster Samples Operator](#)

3.2.16. Cluster Image Registry capability

Purpose

The Cluster Image Registry Operator provides features for the **ImageRegistry** capability.

The Cluster Image Registry Operator manages a singleton instance of the OpenShift image registry. It manages all configuration of the registry, including creating storage.

On initial start up, the Operator creates a default **image-registry** resource instance based on the configuration detected in the cluster. This indicates what cloud storage type to use based on the cloud provider.

If insufficient information is available to define a complete **image-registry** resource, then an incomplete resource is defined and the Operator updates the resource status with information about what is missing.

The Cluster Image Registry Operator runs in the **openshift-image-registry** namespace and it also manages the registry instance in that location. All configuration and workload resources for the registry reside in that namespace.

In order to integrate the image registry into the cluster's user authentication and authorization system, an image pull secret is generated for each service account in the cluster.



IMPORTANT

If you disable the **ImageRegistry** capability or if you disable the integrated OpenShift image registry in the Cluster Image Registry Operator's configuration, the image pull secret is not generated for each service account.

If you disable the **ImageRegistry** capability, you can reduce the overall resource footprint of OpenShift Container Platform in Telco environments. Depending on your deployment, you can disable this component if you do not need it.

Project

[cluster-image-registry-operator](#)

Additional resources

- [Image Registry Operator in OpenShift Container Platform](#)
- [Automatically generated secrets](#)

3.3. ADDITIONAL RESOURCES

- [Enabling cluster capabilities after installation](#)

CHAPTER 4. DISCONNECTED INSTALLATION MIRRORING

4.1. ABOUT DISCONNECTED INSTALLATION MIRRORING

You can use a mirror registry to ensure that your clusters only use container images that satisfy your organizational controls on external content. Before you install a cluster on infrastructure that you provision in a restricted network, you must mirror the required container images into that environment. To mirror container images, you must have a registry for mirroring.

4.1.1. Creating a mirror registry

If you already have a container image registry, such as Red Hat Quay, you can use it as your mirror registry. If you do not already have a registry, you can [create a mirror registry using the *mirror registry for Red Hat OpenShift*](#).

4.1.2. Mirroring images for a disconnected installation

You can use one of the following procedures to mirror your OpenShift Container Platform image repository to your mirror registry:

- [Mirroring images for a disconnected installation](#)
- [Mirroring images for a disconnected installation using the *oc-mirror* plugin](#)


4.2. CREATING A MIRROR REGISTRY WITH MIRROR REGISTRY FOR RED HAT OPENSIFT

The *mirror registry for Red Hat OpenShift* is a small and streamlined container registry that you can use as a target for mirroring the required container images of OpenShift Container Platform for disconnected installations.

If you already have a container image registry, such as Red Hat Quay, you can skip this section and go straight to [Mirroring the OpenShift Container Platform image repository](#).

4.2.1. Prerequisites

- An OpenShift Container Platform subscription.
- Red Hat Enterprise Linux (RHEL) 8 and 9 with Podman 3.4.2 or later and OpenSSL installed.
- Fully qualified domain name for the Red Hat Quay service, which must resolve through a DNS server.
- Key-based SSH connectivity on the target host. SSH keys are automatically generated for local installs. For remote hosts, you must generate your own SSH keys.
- 2 or more vCPUs.
- 8 GB of RAM.
- About 12 GB for OpenShift Container Platform 4.16 release images, or about 358 GB for OpenShift Container Platform 4.16 release images and OpenShift Container Platform 4.16 Red Hat Operator images. Up to 1 TB per stream or more is suggested.



IMPORTANT

These requirements are based on local testing results with only release images and Operator images. Storage requirements can vary based on your organization's needs. You might require more space, for example, when you mirror multiple z-streams. You can use standard [Red Hat Quay functionality](#) or the proper [API callout](#) to remove unnecessary images and free up space.

4.2.2. Mirror registry for Red Hat OpenShift introduction

For disconnected deployments of OpenShift Container Platform, a container registry is required to carry out the installation of the clusters. To run a production-grade registry service on such a cluster, you must create a separate registry deployment to install the first cluster. The *mirror registry for Red Hat OpenShift* addresses this need and is included in every OpenShift subscription. It is available for download on the [OpenShift console Downloads](#) page.

The *mirror registry for Red Hat OpenShift* allows users to install a small-scale version of Red Hat Quay and its required components using the **mirror-registry** command line interface (CLI) tool. The *mirror registry for Red Hat OpenShift* is deployed automatically with preconfigured local storage and a local database. It also includes auto-generated user credentials and access permissions with a single set of inputs and no additional configuration choices to get started.

The *mirror registry for Red Hat OpenShift* provides a pre-determined network configuration and reports deployed component credentials and access URLs upon success. A limited set of optional configuration inputs like fully qualified domain name (FQDN) services, superuser name and password, and custom TLS certificates are also provided. This provides users with a container registry so that they can easily create an offline mirror of all OpenShift Container Platform release content when running OpenShift Container Platform in restricted network environments.

Use of the *mirror registry for Red Hat OpenShift* is optional if another container registry is already available in the install environment.

4.2.2.1. Mirror registry for Red Hat OpenShift limitations

The following limitations apply to the *mirror registry for Red Hat OpenShift*:

- The *mirror registry for Red Hat OpenShift* is not a highly-available registry and only local file system storage is supported. It is not intended to replace Red Hat Quay or the internal image registry for OpenShift Container Platform.
- The *mirror registry for Red Hat OpenShift* is only supported for hosting images that are required to install a disconnected OpenShift Container Platform cluster, such as Release images or Red Hat Operator images. It uses local storage on your Red Hat Enterprise Linux (RHEL) machine, and storage supported by RHEL is supported by the *mirror registry for Red Hat OpenShift*.



NOTE

Because the *mirror registry for Red Hat OpenShift* uses local storage, you should remain aware of the storage usage consumed when mirroring images and use Red Hat Quay's garbage collection feature to mitigate potential issues. For more information about this feature, see "Red Hat Quay garbage collection".

- Support for Red Hat product images that are pushed to the *mirror registry for Red Hat OpenShift* for bootstrapping purposes are covered by valid subscriptions for each respective product. A list of exceptions to further enable the bootstrap experience can be found on the

Self-managed Red Hat OpenShift sizing and subscription guide .

- Content built by customers should not be hosted by the *mirror registry for Red Hat OpenShift* .
- Using the *mirror registry for Red Hat OpenShift* with more than one cluster is discouraged because multiple clusters can create a single point of failure when updating your cluster fleet. It is advised to leverage the *mirror registry for Red Hat OpenShift* to install a cluster that can host a production-grade, highly-available registry such as Red Hat Quay, which can serve OpenShift Container Platform content to other clusters.

4.2.3. Mirroring on a local host with mirror registry for Red Hat OpenShift

This procedure explains how to install the *mirror registry for Red Hat OpenShift* on a local host using the **mirror-registry** installer tool. By doing so, users can create a local host registry running on port 443 for the purpose of storing a mirror of OpenShift Container Platform images.



NOTE

Installing the *mirror registry for Red Hat OpenShift* using the **mirror-registry** CLI tool makes several changes to your machine. After installation, a **\$HOME/quay-install** directory is created, which has installation files, local storage, and the configuration bundle. Trusted SSH keys are generated in case the deployment target is the local host, and systemd files on the host machine are set up to ensure that container runtimes are persistent. Additionally, an initial user named **init** is created with an automatically generated password. All access credentials are printed at the end of the install routine.

Procedure

1. Download the **mirror-registry.tar.gz** package for the latest version of the *mirror registry for Red Hat OpenShift* found on the [OpenShift console Downloads](#) page.
2. Install the *mirror registry for Red Hat OpenShift* on your local host with your current user account by using the **mirror-registry** tool. For a full list of available flags, see "mirror registry for Red Hat OpenShift flags".

```
$ ./mirror-registry install \
  --quayHostname <host_example_com> \
  --quayRoot <example_directory_name>
```

3. Use the user name and password generated during installation to log into the registry by running the following command:

```
$ podman login -u init \
  -p <password> \
  <host_example_com>:8443> \
  --tls-verify=false 1
```

- 1** You can avoid running **--tls-verify=false** by configuring your system to trust the generated rootCA certificates. See "Using SSL to protect connections to Red Hat Quay" and "Configuring the system to trust the certificate authority" for more information.

**NOTE**

You can also log in by accessing the UI at <https://<host.example.com>:8443> after installation.


4. You can mirror OpenShift Container Platform images after logging in. Depending on your needs, see either the "Mirroring the OpenShift Container Platform image repository" or the "Mirroring Operator catalogs for use with disconnected clusters" sections of this document.

**NOTE**

If there are issues with images stored by the *mirror registry for Red Hat OpenShift* due to storage layer problems, you can remirror the OpenShift Container Platform images, or reinstall mirror registry on more stable storage.

4.2.4. Updating mirror registry for Red Hat OpenShift from a local host

This procedure explains how to update the *mirror registry for Red Hat OpenShift* from a local host using the **upgrade** command. Updating to the latest version ensures new features, bug fixes, and security vulnerability fixes.

**IMPORTANT**

When updating, there is intermittent downtime of your mirror registry, as it is restarted during the update process.


Prerequisites

- You have installed the *mirror registry for Red Hat OpenShift* on a local host.

Procedure

- If you are upgrading the *mirror registry for Red Hat OpenShift* from 1.2.z → 1.3.0, and your installation directory is the default at **/etc/quay-install**, you can enter the following command:

```
$ sudo ./mirror-registry upgrade -v
```

**NOTE**

- *mirror registry for Red Hat OpenShift* migrates Podman volumes for Quay storage, Postgres data, and **/etc/quay-install** data to the new **\$HOME/quay-install** location. This allows you to use *mirror registry for Red Hat OpenShift* without the **--quayRoot** flag during future upgrades.
- Users who upgrade *mirror registry for Red Hat OpenShift* with the **./mirror-registry upgrade -v** flag must include the same credentials used when creating their mirror registry. For example, if you installed the *mirror registry for Red Hat OpenShift* with **--quayHostname <host_example_com>** and **--quayRoot <example_directory_name>**, you must include that string to properly upgrade the mirror registry.

- If you are upgrading the *mirror registry for Red Hat OpenShift* from 1.2.z → 1.3.0 and you used a specified directory in your 1.2.z deployment, you must pass in the new **--pgStorage** and **--quayStorage** flags. For example:

```
$ sudo ./mirror-registry upgrade --quayHostname <host_example_com> --quayRoot
<example_directory_name> --pgStorage <example_directory_name>/pg-data --quayStorage
<example_directory_name>/quay-storage -v
```

4.2.5. Mirroring on a remote host with mirror registry for Red Hat OpenShift

This procedure explains how to install the *mirror registry for Red Hat OpenShift* on a remote host using the **mirror-registry** tool. By doing so, users can create a registry to hold a mirror of OpenShift Container Platform images.



NOTE

Installing the *mirror registry for Red Hat OpenShift* using the **mirror-registry** CLI tool makes several changes to your machine. After installation, a **\$HOME/quay-install** directory is created, which has installation files, local storage, and the configuration bundle. Trusted SSH keys are generated in case the deployment target is the local host, and systemd files on the host machine are set up to ensure that container runtimes are persistent. Additionally, an initial user named **init** is created with an automatically generated password. All access credentials are printed at the end of the install routine.

Procedure

1. Download the **mirror-registry.tar.gz** package for the latest version of the *mirror registry for Red Hat OpenShift* found on the [OpenShift console Downloads](#) page.
2. Install the *mirror registry for Red Hat OpenShift* on your local host with your current user account by using the **mirror-registry** tool. For a full list of available flags, see "mirror registry for Red Hat OpenShift flags".

```
$ ./mirror-registry install -v \
--targetHostname <host_example_com> \
--targetUsername <example_user> \
-k ~/.ssh/my_ssh_key \
--quayHostname <host_example_com> \
--quayRoot <example_directory_name>
```

3. Use the user name and password generated during installation to log into the mirror registry by running the following command:

```
$ podman login -u init \
-p <password> \
<host_example_com>:8443 \
--tls-verify=false 1
```

- 1** You can avoid running **--tls-verify=false** by configuring your system to trust the generated rootCA certificates. See "Using SSL to protect connections to Red Hat Quay" and "Configuring the system to trust the certificate authority" for more information.

**NOTE**

You can also log in by accessing the UI at <https://<host.example.com>:8443> after installation.

4. You can mirror OpenShift Container Platform images after logging in. Depending on your needs, see either the "Mirroring the OpenShift Container Platform image repository" or the "Mirroring Operator catalogs for use with disconnected clusters" sections of this document.

**NOTE**

If there are issues with images stored by the *mirror registry for Red Hat OpenShift* due to storage layer problems, you can remirror the OpenShift Container Platform images, or reinstall mirror registry on more stable storage.

4.2.6. Updating mirror registry for Red Hat OpenShift from a remote host

This procedure explains how to update the *mirror registry for Red Hat OpenShift* from a remote host using the **upgrade** command. Updating to the latest version ensures bug fixes and security vulnerability fixes.

**IMPORTANT**

When updating, there is intermittent downtime of your mirror registry, as it is restarted during the update process.

Prerequisites

- You have installed the *mirror registry for Red Hat OpenShift* on a remote host.

Procedure

- To upgrade the *mirror registry for Red Hat OpenShift* from a remote host, enter the following command:

```
$ ./mirror-registry upgrade -v --targetHostname <remote_host_url> --targetUsername
<user_name> -k ~/.ssh/my_ssh_key
```

**NOTE**

Users who upgrade the *mirror registry for Red Hat OpenShift* with the **./mirror-registry upgrade -v** flag must include the same credentials used when creating their mirror registry. For example, if you installed the *mirror registry for Red Hat OpenShift* with **--quayHostname <host_example_com>** and **--quayRoot <example_directory_name>**, you must include that string to properly upgrade the mirror registry.

4.2.7. Replacing mirror registry for Red Hat OpenShift SSL/TLS certificates

In some cases, you might want to update your SSL/TLS certificates for the *mirror registry for Red Hat OpenShift*. This is useful in the following scenarios:

- If you are replacing the current *mirror registry for Red Hat OpenShift* certificate.

- If you are using the same certificate as the previous *mirror registry for Red Hat OpenShift* installation.
- If you are periodically updating the *mirror registry for Red Hat OpenShift* certificate.

Use the following procedure to replace *mirror registry for Red Hat OpenShift* SSL/TLS certificates.

Prerequisites

- You have downloaded the `./mirror-registry` binary from the [OpenShift console Downloads](#) page.

Procedure

1. Enter the following command to install the *mirror registry for Red Hat OpenShift*:

```
$ ./mirror-registry install \
--quayHostname <host_example_com> \
--quayRoot <example_directory_name>
```

This installs the *mirror registry for Red Hat OpenShift* to the `$HOME/quay-install` directory.

2. Prepare a new certificate authority (CA) bundle and generate new `ssl.key` and `ssl.crt` key files. For more information, see [Using SSL/TLS to protect connections to Red Hat Quay](#).
3. Assign `/$HOME/quay-install` an environment variable, for example, `QUAY`, by entering the following command:

```
$ export QUAY=/$HOME/quay-install
```

4. Copy the new `ssl.crt` file to the `/$HOME/quay-install` directory by entering the following command:

```
$ cp ~/ssl.crt $QUAY/quay-config
```

5. Copy the new `ssl.key` file to the `/$HOME/quay-install` directory by entering the following command:

```
$ cp ~/ssl.key $QUAY/quay-config
```

6. Restart the `quay-app` application pod by entering the following command:

```
$ systemctl restart quay-app
```

4.2.8. Uninstalling the mirror registry for Red Hat OpenShift

- You can uninstall the *mirror registry for Red Hat OpenShift* from your local host by running the following command:

```
$ ./mirror-registry uninstall -v \
--quayRoot <example_directory_name>
```

**NOTE**

- Deleting the *mirror registry for Red Hat OpenShift* will prompt the user before deletion. You can use **--autoApprove** to skip this prompt.
- Users who install the *mirror registry for Red Hat OpenShift* with the **--quayRoot** flag must include the **--quayRoot** flag when uninstalling. For example, if you installed the *mirror registry for Red Hat OpenShift* with **--quayRoot example_directory_name**, you must include that string to properly uninstall the mirror registry.

4.2.9. Mirror registry for Red Hat OpenShift flags

The following flags are available for the *mirror registry for Red Hat OpenShift*:

Flags	Description
--autoApprove	A boolean value that disables interactive prompts. If set to true , the quayRoot directory is automatically deleted when uninstalling the mirror registry. Defaults to false if left unspecified.
--initPassword	The password of the init user created during Quay installation. Must be at least eight characters and contain no whitespace.
--initUser string	Shows the username of the initial user. Defaults to init if left unspecified.
--no-color, -c	Allows users to disable color sequences and propagate that to Ansible when running install, uninstall, and upgrade commands.
--pgStorage	The folder where Postgres persistent storage data is saved. Defaults to the pg-storage Podman volume. Root privileges are required to uninstall.
--quayHostname	The fully-qualified domain name of the mirror registry that clients will use to contact the registry. Equivalent to SERVER_HOSTNAME in the Quay config.yaml . Must resolve by DNS. Defaults to <targetHostname>:8443 if left unspecified. ^[1]
--quayStorage	The folder where Quay persistent storage data is saved. Defaults to the quay-storage Podman volume. Root privileges are required to uninstall.
--quayRoot, -r	The directory where container image layer and configuration data is saved, including rootCA.key , rootCA.pem , and rootCA.srl certificates. Defaults to \$HOME/quay-install if left unspecified.
--ssh-key, -k	The path of your SSH identity key. Defaults to ~/ssh/quay_installer if left unspecified.
--sslCert	The path to the SSL/TLS public key / certificate. Defaults to {quayRoot}/quay-config and is auto-generated if left unspecified.

Flags	Description
--sslCheckSkip	Skips the check for the certificate hostname against the SERVER_HOSTNAME in the config.yaml file. ^[2]
--sslKey	The path to the SSL/TLS private key used for HTTPS communication. Defaults to {quayRoot}/quay-config and is auto-generated if left unspecified.
--targetHostname, -H	The hostname of the target you want to install Quay to. Defaults to \$HOST , for example, a local host, if left unspecified.
--targetUsername, -u	The user on the target host which will be used for SSH. Defaults to \$USER , for example, the current user if left unspecified.
--verbose, -v	Shows debug logs and Ansible playbook outputs.
--version	Shows the version for the <i>mirror registry for Red Hat OpenShift</i>

1. **--quayHostname** must be modified if the public DNS name of your system is different from the local hostname. Additionally, the **--quayHostname** flag does not support installation with an IP address. Installation with a hostname is required.
2. **--sslCheckSkip** is used in cases when the mirror registry is set behind a proxy and the exposed hostname is different from the internal Quay hostname. It can also be used when users do not want the certificates to be validated against the provided Quay hostname during installation.

4.2.10. Mirror registry for Red Hat OpenShift release notes

The *mirror registry for Red Hat OpenShift* is a small and streamlined container registry that you can use as a target for mirroring the required container images of OpenShift Container Platform for disconnected installations.

These release notes track the development of the *mirror registry for Red Hat OpenShift* in OpenShift Container Platform.

For an overview of the *mirror registry for Red Hat OpenShift*, see [Creating a mirror registry with mirror registry for Red Hat OpenShift](#).

4.2.10.1. Mirror registry for Red Hat OpenShift 1.3.11

Issued: 2024-04-23

Mirror registry for Red Hat OpenShift is now available with Red Hat Quay 3.8.15.

The following advisory is available for the *mirror registry for Red Hat OpenShift*:

- [RHBA-2024:1758 - mirror registry for Red Hat OpenShift 1.3.11](#)

4.2.10.2. Mirror registry for Red Hat OpenShift 1.3.10

Issued: 2023-12-07

Mirror registry for Red Hat OpenShift is now available with Red Hat Quay 3.8.14.

The following advisory is available for the *mirror registry for Red Hat OpenShift*:

- [RHBA-2023:7628 - mirror registry for Red Hat OpenShift 1.3.10](#)

4.2.10.3. Mirror registry for Red Hat OpenShift 1.3.9

Issued: 2023-09-19

Mirror registry for Red Hat OpenShift is now available with Red Hat Quay 3.8.12.

The following advisory is available for the *mirror registry for Red Hat OpenShift*:

- [RHBA-2023:5241 - mirror registry for Red Hat OpenShift 1.3.9](#)

4.2.10.4. Mirror registry for Red Hat OpenShift 1.3.8

Issued: 2023-08-16

Mirror registry for Red Hat OpenShift is now available with Red Hat Quay 3.8.11.

The following advisory is available for the *mirror registry for Red Hat OpenShift*:

- [RHBA-2023:4622 - mirror registry for Red Hat OpenShift 1.3.8](#)

4.2.10.5. Mirror registry for Red Hat OpenShift 1.3.7

Issued: 2023-07-19

Mirror registry for Red Hat OpenShift is now available with Red Hat Quay 3.8.10.

The following advisory is available for the *mirror registry for Red Hat OpenShift*:

- [RHBA-2023:4087 - mirror registry for Red Hat OpenShift 1.3.7](#)

4.2.10.6. Mirror registry for Red Hat OpenShift 1.3.6

Issued: 2023-05-30

Mirror registry for Red Hat OpenShift is now available with Red Hat Quay 3.8.8.

The following advisory is available for the *mirror registry for Red Hat OpenShift*:

- [RHBA-2023:3302 - mirror registry for Red Hat OpenShift 1.3.6](#)

4.2.10.7. Mirror registry for Red Hat OpenShift 1.3.5

Issued: 2023-05-18

Mirror registry for Red Hat OpenShift is now available with Red Hat Quay 3.8.7.

The following advisory is available for the *mirror registry for Red Hat OpenShift*:

- [RHBA-2023:3225 - mirror registry for Red Hat OpenShift 1.3.5](#)

4.2.10.8. Mirror registry for Red Hat OpenShift 1.3.4

Issued: 2023-04-25

Mirror registry for Red Hat OpenShift is now available with Red Hat Quay 3.8.6.

The following advisory is available for the *mirror registry for Red Hat OpenShift*:

- [RHBA-2023:1914 - mirror registry for Red Hat OpenShift 1.3.4](#)

4.2.10.9. Mirror registry for Red Hat OpenShift 1.3.3

Issued: 2023-04-05

Mirror registry for Red Hat OpenShift is now available with Red Hat Quay 3.8.5.

The following advisory is available for the *mirror registry for Red Hat OpenShift*:

- [RHBA-2023:1528 - mirror registry for Red Hat OpenShift 1.3.3](#)

4.2.10.10. Mirror registry for Red Hat OpenShift 1.3.2

Issued: 2023-03-21

Mirror registry for Red Hat OpenShift is now available with Red Hat Quay 3.8.4.

The following advisory is available for the *mirror registry for Red Hat OpenShift*:

- [RHBA-2023:1376 - mirror registry for Red Hat OpenShift 1.3.2](#)

4.2.10.11. Mirror registry for Red Hat OpenShift 1.3.1

Issued: 2023-03-7

Mirror registry for Red Hat OpenShift is now available with Red Hat Quay 3.8.3.

The following advisory is available for the *mirror registry for Red Hat OpenShift*:

- [RHBA-2023:1086 - mirror registry for Red Hat OpenShift 1.3.1](#)

4.2.10.12. Mirror registry for Red Hat OpenShift 1.3.0

Issued: 2023-02-20

Mirror registry for Red Hat OpenShift is now available with Red Hat Quay 3.8.1.

The following advisory is available for the *mirror registry for Red Hat OpenShift*:

- [RHBA-2023:0558 - mirror registry for Red Hat OpenShift 1.3.0](#)

4.2.10.12.1. New features

- *Mirror registry for Red Hat OpenShift* is now supported on Red Hat Enterprise Linux (RHEL) 9 installations.

- IPv6 support is now available on *mirror registry for Red Hat OpenShift* local host installations. IPv6 is currently unsupported on *mirror registry for Red Hat OpenShift* remote host installations.
- A new feature flag, **--quayStorage**, has been added. By specifying this flag, you can manually set the location for the Quay persistent storage.
- A new feature flag, **--pgStorage**, has been added. By specifying this flag, you can manually set the location for the Postgres persistent storage.
- Previously, users were required to have root privileges (**sudo**) to install *mirror registry for Red Hat OpenShift*. With this update, **sudo** is no longer required to install *mirror registry for Red Hat OpenShift*.

When *mirror registry for Red Hat OpenShift* was installed with **sudo**, an `/etc/quay-install` directory that contained installation files, local storage, and the configuration bundle was created. With the removal of the **sudo** requirement, installation files and the configuration bundle are now installed to `$HOME/quay-install`. Local storage, for example Postgres and Quay, are now stored in named volumes automatically created by Podman.

To override the default directories that these files are stored in, you can use the command line arguments for *mirror registry for Red Hat OpenShift*. For more information about *mirror registry for Red Hat OpenShift* command line arguments, see "[Mirror registry for Red Hat OpenShift flags](#)".

4.2.10.12.2. Bug fixes

- Previously, the following error could be returned when attempting to uninstall *mirror registry for Red Hat OpenShift*: **["Error: no container with name or ID \"quay-postgres\" found: no such container"], "stdout": "", "stdout_lines": []***. With this update, the order that *mirror registry for Red Hat OpenShift* services are stopped and uninstalled have been changed so that the error no longer occurs when uninstalling *mirror registry for Red Hat OpenShift*. For more information, see [PROJQUAY-4629](#).

4.2.10.13. Mirror registry for Red Hat OpenShift 1.2.9

Mirror registry for Red Hat OpenShift is now available with Red Hat Quay 3.7.10.

The following advisory is available for the *mirror registry for Red Hat OpenShift*:

- [RHBA-2022:7369 - mirror registry for Red Hat OpenShift 1.2.9](#)

4.2.10.14. Mirror registry for Red Hat OpenShift 1.2.8

Mirror registry for Red Hat OpenShift is now available with Red Hat Quay 3.7.9.

The following advisory is available for the *mirror registry for Red Hat OpenShift*:

- [RHBA-2022:7065 - mirror registry for Red Hat OpenShift 1.2.8](#)

4.2.10.15. Mirror registry for Red Hat OpenShift 1.2.7

Mirror registry for Red Hat OpenShift is now available with Red Hat Quay 3.7.8.

The following advisory is available for the *mirror registry for Red Hat OpenShift*:

- [RHBA-2022:6500 - mirror registry for Red Hat OpenShift 1.2.7](#)

4.2.10.15.1. Bug fixes

- Previously, **getFQDN()** relied on the fully-qualified domain name (FQDN) library to determine its FQDN, and the FQDN library tried to read the **/etc/hosts** folder directly. Consequently, on some Red Hat Enterprise Linux CoreOS (RHCOS) installations with uncommon DNS configurations, the FQDN library would fail to install and abort the installation. With this update, *mirror registry for Red Hat OpenShift* uses **hostname** to determine the FQDN. As a result, the FQDN library does not fail to install. ([PROJQUAY-4139](#))

4.2.10.16. Mirror registry for Red Hat OpenShift 1.2.6

Mirror registry for Red Hat OpenShift is now available with Red Hat Quay 3.7.7.

The following advisory is available for the *mirror registry for Red Hat OpenShift*:

- [RHBA-2022:6278 - mirror registry for Red Hat OpenShift 1.2.6](#)

4.2.10.16.1. New features

A new feature flag, **--no-color (-c)** has been added. This feature flag allows users to disable color sequences and propagate that to Ansible when running install, uninstall, and upgrade commands.

4.2.10.17. Mirror registry for Red Hat OpenShift 1.2.5

Mirror registry for Red Hat OpenShift is now available with Red Hat Quay 3.7.6.

The following advisory is available for the *mirror registry for Red Hat OpenShift*:

- [RHBA-2022:6071 - mirror registry for Red Hat OpenShift 1.2.5](#)

4.2.10.18. Mirror registry for Red Hat OpenShift 1.2.4

Mirror registry for Red Hat OpenShift is now available with Red Hat Quay 3.7.5.

The following advisory is available for the *mirror registry for Red Hat OpenShift*:

- [RHBA-2022:5884 - mirror registry for Red Hat OpenShift 1.2.4](#)

4.2.10.19. Mirror registry for Red Hat OpenShift 1.2.3

Mirror registry for Red Hat OpenShift is now available with Red Hat Quay 3.7.4.

The following advisory is available for the *mirror registry for Red Hat OpenShift*:

- [RHBA-2022:5649 - mirror registry for Red Hat OpenShift 1.2.3](#)

4.2.10.20. Mirror registry for Red Hat OpenShift 1.2.2

Mirror registry for Red Hat OpenShift is now available with Red Hat Quay 3.7.3.

The following advisory is available for the *mirror registry for Red Hat OpenShift*:

- [RHBA-2022:5501 - mirror registry for Red Hat OpenShift 1.2.2](#)

4.2.10.21. Mirror registry for Red Hat OpenShift 1.2.1

Mirror registry for Red Hat OpenShift is now available with Red Hat Quay 3.7.2.

The following advisory is available for the *mirror registry for Red Hat OpenShift*:

- [RHBA-2022:4986 - mirror registry for Red Hat OpenShift 1.2.1](#)

4.2.10.22. Mirror registry for Red Hat OpenShift 1.2.0

Mirror registry for Red Hat OpenShift is now available with Red Hat Quay 3.7.1.

The following advisory is available for the *mirror registry for Red Hat OpenShift*:

- [RHBA-2022:4986 - mirror registry for Red Hat OpenShift 1.2.0](#)

4.2.10.22.1. Bug fixes

- Previously, all components and workers running inside of the Quay pod Operator had log levels set to **DEBUG**. As a result, large traffic logs were created that consumed unnecessary space. With this update, log levels are set to **WARN** by default, which reduces traffic information while emphasizing problem scenarios. ([PROJQUAY-3504](#))


4.2.10.23. Mirror registry for Red Hat OpenShift 1.1.0

The following advisory is available for the *mirror registry for Red Hat OpenShift*:

- [RHBA-2022:0956 - mirror registry for Red Hat OpenShift 1.1.0](#)

4.2.10.23.1. New features

- A new command, **mirror-registry upgrade** has been added. This command upgrades all container images without interfering with configurations or data.



NOTE

If **quayRoot** was previously set to something other than default, it must be passed into the upgrade command.

4.2.10.23.2. Bug fixes

- Previously, the absence of **quayHostname** or **targetHostname** did not default to the local hostname. With this update, **quayHostname** and **targetHostname** now default to the local hostname if they are missing. ([PROJQUAY-3079](#))
- Previously, the command `./mirror-registry --version` returned an **unknown flag** error. Now, running `./mirror-registry --version` returns the current version of the *mirror registry for Red Hat OpenShift*. ([PROJQUAY-3086](#))
- Previously, users could not set a password during installation, for example, when running `./mirror-registry install --initUser <user_name> --initPassword <password> --verbose`. With this update, users can set a password during installation. ([PROJQUAY-3149](#))
- Previously, the *mirror registry for Red Hat OpenShift* did not recreate pods if they were destroyed. Now, pods are recreated if they are destroyed. ([PROJQUAY-3261](#))

4.2.11. Troubleshooting mirror registry for Red Hat OpenShift

To assist in troubleshooting *mirror registry for Red Hat OpenShift*, you can gather logs of systemd services installed by the mirror registry. The following services are installed:

- quay-app.service
- quay-postgres.service
- quay-redis.service
- quay-pod.service

Prerequisites

- You have installed *mirror registry for Red Hat OpenShift*.

Procedure

- If you installed *mirror registry for Red Hat OpenShift* with root privileges, you can get the status information of its systemd services by entering the following command:

```
$ sudo systemctl status <service>
```

- If you installed *mirror registry for Red Hat OpenShift* as a standard user, you can get the status information of its systemd services by entering the following command:

```
$ systemctl --user status <service>
```

4.2.12. Additional resources

- [Red Hat Quay garbage collection](#)
- [Using SSL to protect connections to Red Hat Quay](#)
- [Configuring the system to trust the certificate authority](#)
- [Mirroring the OpenShift Container Platform image repository](#)
- [Mirroring Operator catalogs for use with disconnected clusters](#)

4.3. MIRRORING IMAGES FOR A DISCONNECTED INSTALLATION

You can ensure your clusters only use container images that satisfy your organizational controls on external content. Before you install a cluster on infrastructure that you provision in a restricted network, you must mirror the required container images into that environment. To mirror container images, you must have a registry for mirroring.



IMPORTANT

You must have access to the internet to obtain the necessary container images. In this procedure, you place your mirror registry on a mirror host that has access to both your network and the internet. If you do not have access to a mirror host, use the [Mirroring Operator catalogs for use with disconnected clusters](#) procedure to copy images to a device you can move across network boundaries with.

4.3.1. Prerequisites

- You must have a container image registry that supports [Docker v2-2](#) in the location that will host the OpenShift Container Platform cluster, such as one of the following registries:
 - [Red Hat Quay](#)
 - [JFrog Artifactory](#)
 - [Sonatype Nexus Repository](#)
 - [Harbor](#)

If you have an entitlement to Red Hat Quay, see the documentation on deploying Red Hat Quay [for proof-of-concept purposes](#) or [by using the Red Hat Quay Operator](#) . If you need additional assistance selecting and installing a registry, contact your sales representative or Red Hat support.

- If you do not already have an existing solution for a container image registry, subscribers of OpenShift Container Platform are provided a [mirror registry for Red Hat OpenShift](#) . The *mirror registry for Red Hat OpenShift* is included with your subscription and is a small-scale container registry that can be used to mirror the required container images of OpenShift Container Platform in disconnected installations.

4.3.2. About the mirror registry

You can mirror the images that are required for OpenShift Container Platform installation and subsequent product updates to a container mirror registry such as Red Hat Quay, JFrog Artifactory, Sonatype Nexus Repository, or Harbor. If you do not have access to a large-scale container registry, you can use the *mirror registry for Red Hat OpenShift* , a small-scale container registry included with OpenShift Container Platform subscriptions.

You can use any container registry that supports [Docker v2-2](#), such as Red Hat Quay, the *mirror registry for Red Hat OpenShift*, Artifactory, Sonatype Nexus Repository, or Harbor. Regardless of your chosen registry, the procedure to mirror content from Red Hat hosted sites on the internet to an isolated image registry is the same. After you mirror the content, you configure each cluster to retrieve this content from your mirror registry.




IMPORTANT

The OpenShift image registry cannot be used as the target registry because it does not support pushing without a tag, which is required during the mirroring process.

If choosing a container registry that is not the *mirror registry for Red Hat OpenShift* , it must be reachable by every machine in the clusters that you provision. If the registry is unreachable, installation, updating, or normal operations such as workload relocation might fail. For that reason, you must run mirror registries in a highly available way, and the mirror registries must at least match the production availability of your OpenShift Container Platform clusters.

When you populate your mirror registry with OpenShift Container Platform images, you can follow two scenarios. If you have a host that can access both the internet and your mirror registry, but not your cluster nodes, you can directly mirror the content from that machine. This process is referred to as *connected mirroring*. If you have no such host, you must mirror the images to a file system and then bring that host or removable media into your restricted environment. This process is referred to as *disconnected mirroring*.

For mirrored registries, to view the source of pulled images, you must review the **Trying to access** log entry in the CRI-O logs. Other methods to view the image pull source, such as using the **crictl images** command on a node, show the non-mirrored image name, even though the image is pulled from the mirrored location.



NOTE

Red Hat does not test third party registries with OpenShift Container Platform.

Additional information


For information about viewing the CRI-O logs to view the image source, see [Viewing the image pull source](#).

4.3.3. Preparing your mirror host

Before you perform the mirror procedure, you must prepare the host to retrieve content and push it to the remote location.

4.3.3.1. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.



IMPORTANT

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**. To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:


```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.



NOTE

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your PATH.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

4.3.4. Configuring credentials that allow images to be mirrored

Create a container image registry credentials file that enables you to mirror images from Red Hat to your mirror.



WARNING

Do not use this image registry credentials file as the pull secret when you install a cluster. If you provide this file when you install cluster, all of the machines in the cluster will have write access to your mirror registry.



WARNING

This process requires that you have write access to a container image registry on the mirror registry and adds the credentials to a registry pull secret.

Prerequisites

- You configured a mirror registry to use in your disconnected environment.
- You identified an image repository location on your mirror registry to mirror images into.
- You provisioned a mirror registry account that allows images to be uploaded to that image repository.

Procedure

Complete the following steps on the installation host:

1. Download your **registry.redhat.io** [pull secret from Red Hat OpenShift Cluster Manager](#) .
2. Make a copy of your pull secret in JSON format:

```
$ cat ./pull-secret | jq . > <path>/<pull_secret_file_in_json> 1
```

- 1** Specify the path to the folder to store the pull secret in and a name for the JSON file that you create.

The contents of the file resemble the following example:

```
{
```

```

"auths": {
  "cloud.openshift.com": {
    "auth": "b3BlbnNo...",
    "email": "you@example.com"
  },
  "quay.io": {
    "auth": "b3BlbnNo...",
    "email": "you@example.com"
  },
  "registry.connect.redhat.com": {
    "auth": "NTE3Njg5Nj...",
    "email": "you@example.com"
  },
  "registry.redhat.io": {
    "auth": "NTE3Njg5Nj...",
    "email": "you@example.com"
  }
}
}
}

```

1. Generate the base64-encoded user name and password or token for your mirror registry:

```

$ echo -n '<user_name>:<password>' | base64 -w0 ❶
BGVtbYk3ZHAAtqXs=

```

- ❶ For **<user_name>** and **<password>**, specify the user name and password that you configured for your registry.

2. Edit the JSON file and add a section that describes your registry to it:

```

"auths": {
  "<mirror_registry>": { ❶
    "auth": "<credentials>", ❷
    "email": "you@example.com"
  }
},

```

- ❶ Specify the registry domain name, and optionally the port, that your mirror registry uses to serve content. For example, **registry.example.com** or **registry.example.com:8443**
- ❷ Specify the base64-encoded user name and password for the mirror registry.

The file resembles the following example:

```

{
  "auths": {
    "registry.example.com": {
      "auth": "BGVtbYk3ZHAAtqXs=",
      "email": "you@example.com"
    },
    "cloud.openshift.com": {
      "auth": "b3BlbnNo...",
      "email": "you@example.com"
    }
  }
}

```

```

    },
    "quay.io": {
      "auth": "b3BlbnNo...",
      "email": "you@example.com"
    },
    "registry.connect.redhat.com": {
      "auth": "NTE3Njg5Nj...",
      "email": "you@example.com"
    },
    "registry.redhat.io": {
      "auth": "NTE3Njg5Nj...",
      "email": "you@example.com"
    }
  }
}
}

```

4.3.5. Mirroring the OpenShift Container Platform image repository

Mirror the OpenShift Container Platform image repository to your registry to use during cluster installation or upgrade.

Prerequisites

- Your mirror host has access to the internet.
- You configured a mirror registry to use in your restricted network and can access the certificate and credentials that you configured.
- You downloaded the [pull secret from Red Hat OpenShift Cluster Manager](#) and modified it to include authentication to your mirror repository.
- If you use self-signed certificates, you have specified a Subject Alternative Name in the certificates.

Procedure

Complete the following steps on the mirror host:

1. Review the [OpenShift Container Platform downloads page](#) to determine the version of OpenShift Container Platform that you want to install and determine the corresponding tag on the [Repository Tags](#) page.
2. Set the required environment variables:
 - a. Export the release version:

```
$ OCP_RELEASE=<release_version>
```

For **<release_version>**, specify the tag that corresponds to the version of OpenShift Container Platform to install, such as **4.5.4**.

- b. Export the local registry name and host port:

```
$ LOCAL_REGISTRY='<local_registry_host_name>:<local_registry_host_port>'
```

For **<local_registry_host_name>**, specify the registry domain name for your mirror repository, and for **<local_registry_host_port>**, specify the port that it serves content on.

- c. Export the local repository name:

```
$ LOCAL_REPOSITORY='<local_repository_name>'
```

For **<local_repository_name>**, specify the name of the repository to create in your registry, such as **ocp4/openshift4**.

- d. Export the name of the repository to mirror:

```
$ PRODUCT_REPO='openshift-release-dev'
```

For a production release, you must specify **openshift-release-dev**.

- e. Export the path to your registry pull secret:

```
$ LOCAL_SECRET_JSON='<path_to_pull_secret>'
```

For **<path_to_pull_secret>**, specify the absolute path to and file name of the pull secret for your mirror registry that you created.

- f. Export the release mirror:

```
$ RELEASE_NAME="ocp-release"
```

For a production release, you must specify **ocp-release**.

- g. Export the type of architecture for your cluster:

```
$ ARCHITECTURE=<cluster_architecture> 1
```

- 1 Specify the architecture of the cluster, such as **x86_64**, **aarch64**, **s390x**, or **ppc64le**.

- h. Export the path to the directory to host the mirrored images:

```
$ REMOVABLE_MEDIA_PATH=<path> 1
```

- 1 Specify the full path, including the initial forward slash (/) character.

3. Mirror the version images to the mirror registry:

- If your mirror host does not have internet access, take the following actions:
 - i. Connect the removable media to a system that is connected to the internet.
 - ii. Review the images and configuration manifests to mirror:

```
$ oc adm release mirror -a ${LOCAL_SECRET_JSON} \
  --from=quay.io/${PRODUCT_REPO}/${RELEASE_NAME}:${OCP_RELEASE}-
  ${ARCHITECTURE} \
  --to=${LOCAL_REGISTRY}/${LOCAL_REPOSITORY} \
```

```
--to-release-
image=${LOCAL_REGISTRY}/${LOCAL_REPOSITORY}:${OCP_RELEASE}-
${ARCHITECTURE} --dry-run
```

- iii. Record the entire **imageContentSources** section from the output of the previous command. The information about your mirrors is unique to your mirrored repository, and you must add the **imageContentSources** section to the **install-config.yaml** file during installation.
- iv. Mirror the images to a directory on the removable media:

```
$ oc adm release mirror -a ${LOCAL_SECRET_JSON} --to-
dir=${REMOVABLE_MEDIA_PATH}/mirror
quay.io/${PRODUCT_REPO}/${RELEASE_NAME}:${OCP_RELEASE}-
${ARCHITECTURE}
```

- v. Take the media to the restricted network environment and upload the images to the local container registry.

```
$ oc image mirror -a ${LOCAL_SECRET_JSON} --from-
dir=${REMOVABLE_MEDIA_PATH}/mirror
"file://openshift/release:${OCP_RELEASE}*"
${LOCAL_REGISTRY}/${LOCAL_REPOSITORY} 1
```

- 1** For **REMOVABLE_MEDIA_PATH**, you must use the same path that you specified when you mirrored the images.



IMPORTANT

Running **oc image mirror** might result in the following error: **error: unable to retrieve source image**. This error occurs when image indexes include references to images that no longer exist on the image registry. Image indexes might retain older references to allow users running those images an upgrade path to newer points on the upgrade graph. As a temporary workaround, you can use the **--skip-missing** option to bypass the error and continue downloading the image index. For more information, see [Service Mesh Operator mirroring failed](#).

- If the local container registry is connected to the mirror host, take the following actions:
 - i. Directly push the release images to the local registry by using following command:

```
$ oc adm release mirror -a ${LOCAL_SECRET_JSON} \
--from=quay.io/${PRODUCT_REPO}/${RELEASE_NAME}:${OCP_RELEASE}-
${ARCHITECTURE} \
--to=${LOCAL_REGISTRY}/${LOCAL_REPOSITORY} \
--to-release-
image=${LOCAL_REGISTRY}/${LOCAL_REPOSITORY}:${OCP_RELEASE}-
${ARCHITECTURE}
```

This command pulls the release information as a digest, and its output includes the **imageContentSources** data that you require when you install your cluster.

- ii. Record the entire **imageContentSources** section from the output of the previous command. The information about your mirrors is unique to your mirrored repository, and you must add the **imageContentSources** section to the **install-config.yaml** file during installation.

**NOTE**

The image name gets patched to Quay.io during the mirroring process, and the podman images will show Quay.io in the registry on the bootstrap virtual machine.

4. To create the installation program that is based on the content that you mirrored, extract it and pin it to the release:

- If your mirror host does not have internet access, run the following command:

```
$ oc adm release extract -a ${LOCAL_SECRET_JSON} --icsp-file=<file> \ --
command=openshift-install
"${LOCAL_REGISTRY}/${LOCAL_REPOSITORY}:${OCP_RELEASE}"
```

- If the local container registry is connected to the mirror host, run the following command:

```
$ oc adm release extract -a ${LOCAL_SECRET_JSON} --command=openshift-install
"${LOCAL_REGISTRY}/${LOCAL_REPOSITORY}:${OCP_RELEASE}-
${ARCHITECTURE}"
```

**IMPORTANT**

To ensure that you use the correct images for the version of OpenShift Container Platform that you selected, you must extract the installation program from the mirrored content.

You must perform this step on a machine with an active internet connection.

5. For clusters using installer-provisioned infrastructure, run the following command:

```
$ openshift-install
```

4.3.6. The Cluster Samples Operator in a disconnected environment

In a disconnected environment, you must take additional steps after you install a cluster to configure the Cluster Samples Operator. Review the following information in preparation.

4.3.6.1. Cluster Samples Operator assistance for mirroring

During installation, OpenShift Container Platform creates a config map named **imagestreamtag-to-image** in the **openshift-cluster-samples-operator** namespace. The **imagestreamtag-to-image** config map contains an entry, the populating image, for each image stream tag.

The format of the key for each entry in the data field in the config map is **<image_stream_name>_<image_stream_tag_name>**.

During a disconnected installation of OpenShift Container Platform, the status of the Cluster Samples Operator is set to **Removed**. If you choose to change it to **Managed**, it installs samples.



NOTE

The use of samples in a network-restricted or discontinued environment may require access to services external to your network. Some example services include: Github, Maven Central, npm, RubyGems, PyPi and others. There might be additional steps to take that allow the cluster samples operators's objects to reach the services they require.

You can use this config map as a reference for which images need to be mirrored for your image streams to import.

- While the Cluster Samples Operator is set to **Removed**, you can create your mirrored registry, or determine which existing mirrored registry you want to use.
- Mirror the samples you want to the mirrored registry using the new config map as your guide.
- Add any of the image streams you did not mirror to the **skippedImagestreams** list of the Cluster Samples Operator configuration object.
- Set **samplesRegistry** of the Cluster Samples Operator configuration object to the mirrored registry.
- Then set the Cluster Samples Operator to **Managed** to install the image streams you have mirrored.

4.3.7. Mirroring Operator catalogs for use with disconnected clusters

You can mirror the Operator contents of a Red Hat-provided catalog, or a custom catalog, into a container image registry using the **oc adm catalog mirror** command. The target registry must support [Docker v2-2](#). For a cluster on a restricted network, this registry can be one that the cluster has network access to, such as a mirror registry created during a restricted network cluster installation.



IMPORTANT

- The OpenShift image registry cannot be used as the target registry because it does not support pushing without a tag, which is required during the mirroring process.
- Running **oc adm catalog mirror** might result in the following error: **error: unable to retrieve source image**. This error occurs when image indexes include references to images that no longer exist on the image registry. Image indexes might retain older references to allow users running those images an upgrade path to newer points on the upgrade graph. As a temporary workaround, you can use the **--skip-missing** option to bypass the error and continue downloading the image index. For more information, see [Service Mesh Operator mirroring failed](#).

The **oc adm catalog mirror** command also automatically mirrors the index image that is specified during the mirroring process, whether it be a Red Hat-provided index image or your own custom-built index image, to the target registry. You can then use the mirrored index image to create a catalog source that allows Operator Lifecycle Manager (OLM) to load the mirrored catalog onto your OpenShift Container Platform cluster.

Additional resources

- [Using Operator Lifecycle Manager on restricted networks](#)

4.3.7.1. Prerequisites

Mirroring Operator catalogs for use with disconnected clusters has the following prerequisites:

- Workstation with unrestricted network access.
- **podman** version 1.9.3 or later.
- If you want to filter, or *prune*, an existing catalog and selectively mirror only a subset of Operators, see the following sections:
 - [Installing the opm CLI](#)
 - [Updating or filtering a file-based catalog image](#)
- If you want to mirror a Red Hat-provided catalog, run the following command on your workstation with unrestricted network access to authenticate with **registry.redhat.io**:

```
$ podman login registry.redhat.io
```

- Access to a mirror registry that supports [Docker v2-2](#).
- On your mirror registry, decide which repository, or namespace, to use for storing mirrored Operator content. For example, you might create an **olm-mirror** repository.
- If your mirror registry does not have internet access, connect removable media to your workstation with unrestricted network access.
- If you are working with private registries, including **registry.redhat.io**, set the **REG_CREDS** environment variable to the file path of your registry credentials for use in later steps. For example, for the **podman** CLI:

```
$ REG_CREDS=${XDG_RUNTIME_DIR}/containers/auth.json
```

4.3.7.2. Extracting and mirroring catalog contents

The **oc adm catalog mirror** command extracts the contents of an index image to generate the manifests required for mirroring. The default behavior of the command generates manifests, then automatically mirrors all of the image content from the index image, as well as the index image itself, to your mirror registry.

Alternatively, if your mirror registry is on a completely disconnected, or *airgapped*, host, you can first mirror the content to removable media, move the media to the disconnected environment, then mirror the content from the media to the registry.

4.3.7.2.1. Mirroring catalog contents to registries on the same network

If your mirror registry is co-located on the same network as your workstation with unrestricted network access, take the following actions on your workstation.

Procedure

1. If your mirror registry requires authentication, run the following command to log in to the registry:

```
$ podman login <mirror_registry>
```

2. Run the following command to extract and mirror the content to the mirror registry:

```
$ oc adm catalog mirror \
  <index_image> \ 1
  <mirror_registry>:<port>[/<repository>] \ 2
  [-a ${REG_CREDS}] \ 3
  [--insecure] \ 4
  [--index-filter-by-os='<platform>/<arch>'] \ 5
  [--manifests-only] 6
```

- 1** Specify the index image for the catalog that you want to mirror.
- 2** Specify the fully qualified domain name (FQDN) for the target registry to mirror the Operator contents to. The mirror registry **<repository>** can be any existing repository, or namespace, on the registry, for example **olm-mirror** as outlined in the prerequisites. If there is an existing repository found during mirroring, the repository name is added to the resulting image name. If you do not want the image name to include the repository name, omit the **<repository>** value from this line, for example **<mirror_registry>:<port>**.
- 3** Optional: If required, specify the location of your registry credentials file. **{REG_CREDS}** is required for **registry.redhat.io**.
- 4** Optional: If you do not want to configure trust for the target registry, add the **--insecure** flag.
- 5** Optional: Specify which platform and architecture of the index image to select when multiple variants are available. Images are passed as '**<platform>/<arch>[/<variant>]**'. This does not apply to images referenced by the index. Valid values are **linux/amd64**, **linux/ppc64le**, **linux/s390x**, **linux/arm64**.
- 6** Optional: Generate only the manifests required for mirroring without actually mirroring the image content to a registry. This option can be useful for reviewing what will be mirrored, and lets you make any changes to the mapping list, if you require only a subset of packages. You can then use the **mapping.txt** file with the **oc image mirror** command to mirror the modified list of images in a later step. This flag is intended for only advanced selective mirroring of content from the catalog.

Example output

```
src image has index label for database path: /database/index.db
using database path mapping: /database/index.db:/tmp/153048078
wrote database to /tmp/153048078 1
...
wrote mirroring manifests to manifests-redhat-operator-index-1614211642 2
```

- 1** Directory for the temporary **index.db** database generated by the command.
- 2**

Record the manifests directory name that is generated. This directory is referenced in subsequent procedures.



NOTE

Red Hat Quay does not support nested repositories. As a result, running the **oc adm catalog mirror** command will fail with a **401** unauthorized error. As a workaround, you can use the **--max-components=2** option when running the **oc adm catalog mirror** command to disable the creation of nested repositories. For more information on this workaround, see the [Unauthorized error thrown while using catalog mirror command with Quay registry](#) Knowledgebase Solution.

Additional resources

- [Architecture and operating system support for Operators](#)

4.3.7.2.2. Mirroring catalog contents to airgapped registries

If your mirror registry is on a completely disconnected, or airgapped, host, take the following actions.

Procedure

1. Run the following command on your workstation with unrestricted network access to mirror the content to local files:

```
$ oc adm catalog mirror \
  <index_image> \ 1
  file:///local/index \ 2
  -a ${REG_CREDS} \ 3
  --insecure \ 4
  --index-filter-by-os='<platform>/<arch>' 5
```

- 1** Specify the index image for the catalog that you want to mirror.
- 2** Specify the content to mirror to local files in your current directory.
- 3** Optional: If required, specify the location of your registry credentials file.
- 4** Optional: If you do not want to configure trust for the target registry, add the **--insecure** flag.
- 5** Optional: Specify which platform and architecture of the index image to select when multiple variants are available. Images are specified as '**<platform>/<arch>/[<variant>]**'. This does not apply to images referenced by the index. Valid values are **linux/amd64**, **linux/ppc64le**, **linux/s390x**, **linux/arm64**, and **.***

Example output

```
...
info: Mirroring completed in 5.93s (5.915MB/s)
wrote mirroring manifests to manifests-my-index-1614985528 1
```

To upload local images to a registry, run:

```
oc adm catalog mirror file:///local/index/myrepo/my-index:v1 REGISTRY/REPOSITORY 2
```

- 1** Record the manifests directory name that is generated. This directory is referenced in subsequent procedures.
- 2** Record the expanded **file://** path that is based on your provided index image. This path is referenced in a subsequent step.

This command creates a **v2/** directory in your current directory.

2. Copy the **v2/** directory to removable media.
3. Physically remove the media and attach it to a host in the disconnected environment that has access to the mirror registry.
4. If your mirror registry requires authentication, run the following command on your host in the disconnected environment to log in to the registry:

```
$ podman login <mirror_registry>
```

5. Run the following command from the parent directory containing the **v2/** directory to upload the images from local files to the mirror registry:

```
$ oc adm catalog mirror \  
  file:///local/index/<repository>/<index_image>:<tag> \1 \  
  <mirror_registry>:<port>[/<repository>] \2 \  
  -a ${REG_CREDS} \3 \  
  --insecure \4 \  
  --index-filter-by-os='<platform>/<arch>' 5
```

- 1** Specify the **file://** path from the previous command output.
- 2** Specify the fully qualified domain name (FQDN) for the target registry to mirror the Operator contents to. The mirror registry **<repository>** can be any existing repository, or namespace, on the registry, for example **olm-mirror** as outlined in the prerequisites. If there is an existing repository found during mirroring, the repository name is added to the resulting image name. If you do not want the image name to include the repository name, omit the **<repository>** value from this line, for example **<mirror_registry>:<port>**.
- 3** Optional: If required, specify the location of your registry credentials file.
- 4** Optional: If you do not want to configure trust for the target registry, add the **--insecure** flag.
- 5** Optional: Specify which platform and architecture of the index image to select when multiple variants are available. Images are specified as **'<platform>/<arch>[/<variant>]'**. This does not apply to images referenced by the index. Valid values are **linux/amd64**, **linux/ppc64le**, **linux/s390x**, **linux/arm64**, and *****.




NOTE

Red Hat Quay does not support nested repositories. As a result, running the **oc adm catalog mirror** command will fail with a **401** unauthorized error. As a workaround, you can use the **--max-components=2** option when running the **oc adm catalog mirror** command to disable the creation of nested repositories. For more information on this workaround, see the [Unauthorized error thrown while using catalog mirror command with Quay registry](#) Knowledgebase Solution.

- Run the **oc adm catalog mirror** command again. Use the newly mirrored index image as the source and the same mirror registry target used in the previous step:

```
$ oc adm catalog mirror \
  <mirror_registry>:<port>/<index_image> \
  <mirror_registry>:<port>/<repository> \
  --manifests-only 1 \
  [-a ${REG_CREDS}] \
  [--insecure]
```

- The **--manifests-only** flag is required for this step so that the command does not copy all of the mirrored content again.



IMPORTANT

This step is required because the image mappings in the **imageContentSourcePolicy.yaml** file generated during the previous step must be updated from local paths to valid mirror locations. Failure to do so will cause errors when you create the **ImageContentSourcePolicy** object in a later step.

After you mirror the catalog, you can continue with the remainder of your cluster installation. After your cluster installation has finished successfully, you must specify the manifests directory from this procedure to create the **ImageContentSourcePolicy** and **CatalogSource** objects. These objects are required to enable installation of Operators from OperatorHub.

Additional resources

- [Architecture and operating system support for Operators](#)

4.3.7.3. Generated manifests

After mirroring Operator catalog content to your mirror registry, a manifests directory is generated in your current directory.

If you mirrored content to a registry on the same network, the directory name takes the following pattern:

```
manifests-<index_image_name>-<random_number>
```

If you mirrored content to a registry on a disconnected host in the previous section, the directory name takes the following pattern:

```
manifests-index/<repository>/<index_image_name>-<random_number>
```

**NOTE**

The manifests directory name is referenced in subsequent procedures.

The manifests directory contains the following files, some of which might require further modification:

- The **catalogSource.yaml** file is a basic definition for a **CatalogSource** object that is pre-populated with your index image tag and other relevant metadata. This file can be used as is or modified to add the catalog source to your cluster.

**IMPORTANT**

If you mirrored the content to local files, you must modify your **catalogSource.yaml** file to remove any backslash (/) characters from the **metadata.name** field. Otherwise, when you attempt to create the object, it fails with an "invalid resource name" error.

- The **imageContentSourcePolicy.yaml** file defines an **ImageContentSourcePolicy** object that can configure nodes to translate between the image references stored in Operator manifests and the mirrored registry.

**NOTE**

If your cluster uses an **ImageContentSourcePolicy** object to configure repository mirroring, you can use only global pull secrets for mirrored registries. You cannot add a pull secret to a project.

- The **mapping.txt** file contains all of the source images and where to map them in the target registry. This file is compatible with the **oc image mirror** command and can be used to further customize the mirroring configuration.

**IMPORTANT**

If you used the **--manifests-only** flag during the mirroring process and want to further trim the subset of packages to mirror, see the steps in the [Mirroring a package manifest format catalog image](#) procedure of the OpenShift Container Platform 4.7 documentation about modifying your **mapping.txt** file and using the file with the **oc image mirror** command.

4.3.7.4. Postinstallation requirements

After you mirror the catalog, you can continue with the remainder of your cluster installation. After your cluster installation has finished successfully, you must specify the manifests directory from this procedure to create the **ImageContentSourcePolicy** and **CatalogSource** objects. These objects are required to populate and enable installation of Operators from OperatorHub.

Additional resources

- [Populating OperatorHub from mirrored Operator catalogs](#)
- [Updating or filtering a file-based catalog image](#)

4.3.8. Next steps

- Install a cluster on infrastructure that you provision in your restricted network, such as on [VMware vSphere](#), [bare metal](#), or [Amazon Web Services](#).

4.3.9. Additional resources

- See [Gathering data about specific features](#) for more information about using must-gather.

4.4. MIRRORING IMAGES FOR A DISCONNECTED INSTALLATION USING THE OC-MIRROR PLUGIN

Running your cluster in a restricted network without direct internet connectivity is possible by installing the cluster from a mirrored set of OpenShift Container Platform container images in a private registry. This registry must be running at all times as long as the cluster is running. See the [Prerequisites](#) section for more information.

You can use the oc-mirror OpenShift CLI (**oc**) plugin to mirror images to a mirror registry in your fully or partially disconnected environments. You must run oc-mirror from a system with internet connectivity in order to download the required images from the official Red Hat registries.

4.4.1. About the oc-mirror plugin

You can use the oc-mirror OpenShift CLI (**oc**) plugin to mirror all required OpenShift Container Platform content and other images to your mirror registry by using a single tool. It provides the following features:

- Provides a centralized method to mirror OpenShift Container Platform releases, Operators, helm charts, and other images.
- Maintains update paths for OpenShift Container Platform and Operators.
- Uses a declarative image set configuration file to include only the OpenShift Container Platform releases, Operators, and images that your cluster needs.
- Performs incremental mirroring, which reduces the size of future image sets.
- Prunes images from the target mirror registry that were excluded from the image set configuration since the previous execution.
- Optionally generates supporting artifacts for OpenShift Update Service (OSUS) usage.

When using the oc-mirror plugin, you specify which content to mirror in an image set configuration file. In this YAML file, you can fine-tune the configuration to only include the OpenShift Container Platform releases and Operators that your cluster needs. This reduces the amount of data that you need to download and transfer. The oc-mirror plugin can also mirror arbitrary helm charts and additional container images to assist users in seamlessly synchronizing their workloads onto mirror registries.

The first time you run the oc-mirror plugin, it populates your mirror registry with the required content to perform your disconnected cluster installation or update. In order for your disconnected cluster to continue receiving updates, you must keep your mirror registry updated. To update your mirror registry, you run the oc-mirror plugin using the same configuration as the first time you ran it. The oc-mirror plugin references the metadata from the storage backend and only downloads what has been released since the last time you ran the tool. This provides update paths for OpenShift Container Platform and Operators and performs dependency resolution as required.

4.4.1.1. High level workflow

The following steps outline the high-level workflow on how to use the oc-mirror plugin to mirror images to a mirror registry:

1. Create an image set configuration file.
2. Mirror the image set to the target mirror registry by using one of the following methods:
 - Mirror an image set directly to the target mirror registry.
 - Mirror an image set to disk, transfer the image set to the target environment, then upload the image set to the target mirror registry.
3. Configure your cluster to use the resources generated by the oc-mirror plugin.
4. Repeat these steps to update your target mirror registry as necessary.



IMPORTANT

When using the oc-mirror CLI plugin to populate a mirror registry, any further updates to the target mirror registry must be made by using the oc-mirror plugin.

4.4.2. oc-mirror plugin compatibility and support

The oc-mirror plugin supports mirroring OpenShift Container Platform payload images and Operator catalogs for OpenShift Container Platform versions 4.12 and later.



NOTE

On **aarch64**, **ppc64le**, and **s390x** architectures the oc-mirror plugin is only supported for OpenShift Container Platform versions 4.14 and later.

Use the latest available version of the oc-mirror plugin regardless of which versions of OpenShift Container Platform you need to mirror.

Additional resources

- For information on updating oc-mirror, see [Viewing the image pull source](#).

4.4.3. About the mirror registry

You can mirror the images that are required for OpenShift Container Platform installation and subsequent product updates to a container mirror registry that supports [Docker v2-2](#), such as Red Hat Quay. If you do not have access to a large-scale container registry, you can use the *mirror registry for Red Hat OpenShift*, which is a small-scale container registry included with OpenShift Container Platform subscriptions.

Regardless of your chosen registry, the procedure to mirror content from Red Hat hosted sites on the internet to an isolated image registry is the same. After you mirror the content, you configure each cluster to retrieve this content from your mirror registry.




IMPORTANT

The OpenShift image registry cannot be used as the target registry because it does not support pushing without a tag, which is required during the mirroring process.

If choosing a container registry that is not the *mirror registry for Red Hat OpenShift*, it must be reachable by every machine in the clusters that you provision. If the registry is unreachable, installation, updating, or normal operations such as workload relocation might fail. For that reason, you must run mirror registries in a highly available way, and the mirror registries must at least match the production availability of your OpenShift Container Platform clusters.

When you populate your mirror registry with OpenShift Container Platform images, you can follow two scenarios. If you have a host that can access both the internet and your mirror registry, but not your cluster nodes, you can directly mirror the content from that machine. This process is referred to as *connected mirroring*. If you have no such host, you must mirror the images to a file system and then bring that host or removable media into your restricted environment. This process is referred to as *disconnected mirroring*.

For mirrored registries, to view the source of pulled images, you must review the **Trying to access** log entry in the CRI-O logs. Other methods to view the image pull source, such as using the **crictl images** command on a node, show the non-mirrored image name, even though the image is pulled from the mirrored location.



NOTE

Red Hat does not test third party registries with OpenShift Container Platform.

Additional resources

- For information about viewing the CRI-O logs to view the image source, see [Viewing the image pull source](#).

4.4.4. Prerequisites

- You must have a container image registry that supports [Docker v2-2](#) in the location that will host the OpenShift Container Platform cluster, such as Red Hat Quay.



NOTE

If you use Red Hat Quay, you must use version 3.6 or later with the `oc-mirror` plugin. If you have an entitlement to Red Hat Quay, see the documentation on deploying Red Hat Quay [for proof-of-concept purposes](#) or [by using the Red Hat Quay Operator](#). If you need additional assistance selecting and installing a registry, contact your sales representative or Red Hat Support.

If you do not already have an existing solution for a container image registry, subscribers of OpenShift Container Platform are provided a [mirror registry for Red Hat OpenShift](#). The *mirror registry for Red Hat OpenShift* is included with your subscription and is a small-scale container registry that can be used to mirror the required container images of OpenShift Container Platform in disconnected installations.

4.4.5. Preparing your mirror hosts

Before you can use the `oc-mirror` plugin to mirror images, you must install the plugin and create a container image registry credentials file to allow the mirroring from Red Hat to your mirror.

4.4.5.1. Installing the `oc-mirror` OpenShift CLI plugin

Install the `oc-mirror` OpenShift CLI plugin to manage image sets in disconnected environments.

Prerequisites

- You have installed the OpenShift CLI (**oc**). If you are mirroring image sets in a fully disconnected environment, ensure the following:
 - You have installed the oc-mirror plugin on the host that has internet access.
 - The host in the disconnected environment has access to the target mirror registry.
- You have set the **umask** parameter to **0022** on the operating system that uses oc-mirror.
- You have installed the correct binary for the RHEL version that you are using.

Procedure


1. Download the oc-mirror CLI plugin.
 - a. Navigate to the [Downloads](#) page of the [OpenShift Cluster Manager](#).
 - b. Under the **OpenShift disconnected installation tools** section, click **Download** for **OpenShift Client (oc) mirror plugin** and save the file.

2. Extract the archive:

```
$ tar xvzf oc-mirror.tar.gz
```

3. If necessary, update the plugin file to be executable:

```
$ chmod +x oc-mirror
```



NOTE

Do not rename the **oc-mirror** file.

4. Install the oc-mirror CLI plugin by placing the file in your **PATH**, for example, **/usr/local/bin**:

```
$ sudo mv oc-mirror /usr/local/bin/.
```

Verification

- Verify that the plugin for oc-mirror v1 is successfully installed by running the following command:

```
$ oc mirror help
```

Additional resources

- [Installing and using CLI plugins](#)

4.4.5.2. Configuring credentials that allow images to be mirrored

Create a container image registry credentials file that enables you to mirror images from Red Hat to your mirror.

**WARNING**

Do not use this image registry credentials file as the pull secret when you install a cluster. If you provide this file when you install cluster, all of the machines in the cluster will have write access to your mirror registry.

**WARNING**

This process requires that you have write access to a container image registry on the mirror registry and adds the credentials to a registry pull secret.

Prerequisites

- You configured a mirror registry to use in your disconnected environment.
- You identified an image repository location on your mirror registry to mirror images into.
- You provisioned a mirror registry account that allows images to be uploaded to that image repository.

Procedure

Complete the following steps on the installation host:

1. Download your **registry.redhat.io** [pull secret from Red Hat OpenShift Cluster Manager](#) .
2. Make a copy of your pull secret in JSON format:

```
$ cat ./pull-secret | jq . > <path>/<pull_secret_file_in_json> 1
```

- 1** Specify the path to the folder to store the pull secret in and a name for the JSON file that you create.

The contents of the file resemble the following example:

```
{
  "auths": {
    "cloud.openshift.com": {
      "auth": "b3BlbnNo...",
      "email": "you@example.com"
    },
    "quay.io": {
      "auth": "b3BlbnNo...",
      "email": "you@example.com"
    },
    "registry.connect.redhat.com": {
```

```

    "auth": "NTE3Njg5Nj...",
    "email": "you@example.com"
  },
  "registry.redhat.io": {
    "auth": "NTE3Njg5Nj...",
    "email": "you@example.com"
  }
}
}
}

```

3. Save the file either as `~/.docker/config.json` or `$XDG_RUNTIME_DIR/containers/auth.json`.

1. Generate the base64-encoded user name and password or token for your mirror registry:

```

$ echo -n '<user_name>:<password>' | base64 -w0 ❶
BGVtbYk3ZHAqXs=

```

❶ For `<user_name>` and `<password>`, specify the user name and password that you configured for your registry.

2. Edit the JSON file and add a section that describes your registry to it:

```

"auths": {
  "<mirror_registry>": { ❶
    "auth": "<credentials>", ❷
    "email": "you@example.com"
  }
},

```

❶ Specify the registry domain name, and optionally the port, that your mirror registry uses to serve content. For example, `registry.example.com` or `registry.example.com:8443`

❷ Specify the base64-encoded user name and password for the mirror registry.

The file resembles the following example:

```

{
  "auths": {
    "registry.example.com": {
      "auth": "BGVtbYk3ZHAqXs=",
      "email": "you@example.com"
    },
    "cloud.openshift.com": {
      "auth": "b3BlbnNo...",
      "email": "you@example.com"
    },
    "quay.io": {
      "auth": "b3BlbnNo...",
      "email": "you@example.com"
    },
    "registry.connect.redhat.com": {
      "auth": "NTE3Njg5Nj...",
      "email": "you@example.com"
    }
  }
}

```

```

    },
    "registry.redhat.io": {
      "auth": "NTE3Njg5Nj...",
      "email": "you@example.com"
    }
  }
}
}

```

4.4.6. Creating the image set configuration

Before you can use the `oc-mirror` plugin to mirror image sets, you must create an image set configuration file. This image set configuration file defines which OpenShift Container Platform releases, Operators, and other images to mirror, along with other configuration settings for the `oc-mirror` plugin.

You must specify a storage backend in the image set configuration file. This storage backend can be a local directory or a registry that supports [Docker v2-2](#). The `oc-mirror` plugin stores metadata in this storage backend during image set creation.



IMPORTANT

Do not delete or modify the metadata that is generated by the `oc-mirror` plugin. You must use the same storage backend every time you run the `oc-mirror` plugin for the same mirror registry.

Prerequisites

- You have created a container image registry credentials file. For instructions, see "Configuring credentials that allow images to be mirrored".

Procedure

- Use the `oc mirror init` command to create a template for the image set configuration and save it to a file called `imageset-config.yaml`:

```
$ oc mirror init <--registry <storage_backend> > imageset-config.yaml 1
```

- Specifies the location of your storage backend, such as `example.com/mirror/oc-mirror-metadata`.

- Edit the file and adjust the settings as necessary:

```

kind: ImageSetConfiguration
apiVersion: mirror.openshift.io/v1alpha2
archiveSize: 4 1
storageConfig: 2
  registry:
    imageURL: example.com/mirror/oc-mirror-metadata 3
    skipTLS: false
  mirror:
    platform:
      channels:
        - name: stable-4.16 4

```

```

type: ocp
graph: true
operators:
- catalog: registry.redhat.io/redhat/redhat-operator-index:v4.16
  packages:
  - name: serverless-operator
  channels:
  - name: stable
additionalImages:
- name: registry.redhat.io/ubi9/ubi:latest
helm: {}

```

- 1 Add **archiveSize** to set the maximum size, in GiB, of each file within the image set.
- 2 Set the back-end location to save the image set metadata to. This location can be a registry or local directory. It is required to specify **storageConfig** values.
- 3 Set the registry URL for the storage backend.
- 4 Set the channel to retrieve the OpenShift Container Platform images from.
- 5 Add **graph: true** to build and push the graph-data image to the mirror registry. The graph-data image is required to create OpenShift Update Service (OSUS). The **graph: true** field also generates the **UpdateService** custom resource manifest. The **oc** command-line interface (CLI) can use the **UpdateService** custom resource manifest to create OSUS. For more information, see *About the OpenShift Update Service*.
- 6 Set the Operator catalog to retrieve the OpenShift Container Platform images from.
- 7 Specify only certain Operator packages to include in the image set. Remove this field to retrieve all packages in the catalog.
- 8 Specify only certain channels of the Operator packages to include in the image set. You must always include the default channel for the Operator package even if you do not use the bundles in that channel. You can find the default channel by running the following command: **oc mirror list operators --catalog=<catalog_name> --package=<package_name>**.
- 9 Specify any additional images to include in image set.



NOTE

The **graph: true** field also mirrors the **ubi-micro** image along with other mirrored images.

See "Image set configuration parameters" for the full list of parameters and "Image set configuration examples" for various mirroring use cases.

3. Save the updated file.
This image set configuration file is required by the **oc mirror** command when mirroring content.

Additional resources

- [Image set configuration parameters](#)

- [Image set configuration examples](#)
- [Using the OpenShift Update Service in a disconnected environment](#)

4.4.7. Mirroring an image set to a mirror registry

You can use the `oc-mirror` CLI plugin to mirror images to a mirror registry in a [partially disconnected environment](#) or in a [fully disconnected environment](#).

These procedures assume that you already have your mirror registry set up.

4.4.7.1. Mirroring an image set in a partially disconnected environment

In a partially disconnected environment, you can mirror an image set directly to the target mirror registry.

4.4.7.1.1. Mirroring from mirror to mirror

You can use the `oc-mirror` plugin to mirror an image set directly to a target mirror registry that is accessible during image set creation.

You are required to specify a storage backend in the image set configuration file. This storage backend can be a local directory or a Docker v2 registry. The `oc-mirror` plugin stores metadata in this storage backend during image set creation.



IMPORTANT

Do not delete or modify the metadata that is generated by the `oc-mirror` plugin. You must use the same storage backend every time you run the `oc-mirror` plugin for the same mirror registry.

Prerequisites

- You have access to the internet to get the necessary container images.
- You have installed the OpenShift CLI (**oc**).
- You have installed the **oc-mirror** CLI plugin.
- You have created the image set configuration file.

Procedure


- Run the **oc mirror** command to mirror the images from the specified image set configuration to a specified registry:

```
$ oc mirror --config=./<imageset-config.yaml> \1
docker://registry.example:5000 \2
```

- 1 Specify the image set configuration file that you created. For example, **imageset-config.yaml**.
- 2 Specify the registry to mirror the image set file to. The registry must start with **docker://**. If you specify a top-level namespace for the mirror registry, you must also use this same namespace on subsequent executions.

Verification

1. Navigate into the **oc-mirror-workspace/** directory that was generated.
2. Navigate into the results directory, for example, **results-1639608409/**.
3. Verify that YAML files are present for the **ImageContentSourcePolicy** and **CatalogSource** resources.



NOTE

The **repositoryDigestMirrors** section of the **ImageContentSourcePolicy** YAML file is used for the **install-config.yaml** file during installation.

Next steps

- Configure your cluster to use the resources generated by oc-mirror.

Troubleshooting

- [Unable to retrieve source image](#) .

4.4.7.2. Mirroring an image set in a fully disconnected environment

To mirror an image set in a fully disconnected environment, you must first [mirror the image set to disk](#) , then [mirror the image set file on disk to a mirror](#) .

4.4.7.2.1. Mirroring from mirror to disk

You can use the oc-mirror plugin to generate an image set and save the contents to disk. The generated image set can then be transferred to the disconnected environment and mirrored to the target registry.



IMPORTANT

Depending on the configuration specified in the image set configuration file, using oc-mirror to mirror images might download several hundreds of gigabytes of data to disk.

The initial image set download when you populate the mirror registry is often the largest. Because you only download the images that changed since the last time you ran the command, when you run the oc-mirror plugin again, the generated image set is often smaller.

You are required to specify a storage backend in the image set configuration file. This storage backend can be a local directory or a docker v2 registry. The oc-mirror plugin stores metadata in this storage backend during image set creation.



IMPORTANT

Do not delete or modify the metadata that is generated by the oc-mirror plugin. You must use the same storage backend every time you run the oc-mirror plugin for the same mirror registry.

Prerequisites

- You have access to the internet to obtain the necessary container images.
- You have installed the OpenShift CLI (**oc**).
- You have installed the **oc-mirror** CLI plugin.
- You have created the image set configuration file.

Procedure

- Run the **oc mirror** command to mirror the images from the specified image set configuration to disk:

```
$ oc mirror --config=./imageset-config.yaml \ 1  
file://<path_to_output_directory> 2
```

- 1 Pass in the image set configuration file that was created. This procedure assumes that it is named **imageset-config.yaml**.
- 2 Specify the target directory where you want to output the image set file. The target directory path must start with **file://**.

Verification

1. Navigate to your output directory:

```
$ cd <path_to_output_directory>
```

2. Verify that an image set **.tar** file was created:

```
$ ls
```

Example output

```
mirror_seq1_000000.tar
```

Next steps

- Transfer the image set **.tar** file to the disconnected environment.

Troubleshooting

- [Unable to retrieve source image](#) .

4.4.7.2.2. Mirroring from disk to mirror

You can use the **oc-mirror** plugin to mirror the contents of a generated image set to the target mirror registry.

Prerequisites

- You have installed the OpenShift CLI (**oc**) in the disconnected environment.

- You have installed the **oc-mirror** CLI plugin in the disconnected environment.
- You have generated the image set file by using the **oc mirror** command.
- You have transferred the image set file to the disconnected environment.

Procedure

- Run the **oc mirror** command to process the image set file on disk and mirror the contents to a target mirror registry:

```
$ oc mirror --from=./mirror_seq1_000000.tar \ 1
docker://registry.example:5000           2
```

- 1 Pass in the image set .tar file to mirror, named **mirror_seq1_000000.tar** in this example. If an **archiveSize** value was specified in the image set configuration file, the image set might be broken up into multiple .tar files. In this situation, you can pass in a directory that contains the image set .tar files.
- 2 Specify the registry to mirror the image set file to. The registry must start with **docker://**. If you specify a top-level namespace for the mirror registry, you must also use this same namespace on subsequent executions.

This command updates the mirror registry with the image set and generates the **ImageContentSourcePolicy** and **CatalogSource** resources.

Verification

1. Navigate into the **oc-mirror-workspace/** directory that was generated.
2. Navigate into the results directory, for example, **results-1639608409/**.
3. Verify that YAML files are present for the **ImageContentSourcePolicy** and **CatalogSource** resources.

Next steps

- Configure your cluster to use the resources generated by oc-mirror.

Troubleshooting

- [Unable to retrieve source image](#) .

4.4.8. Configuring your cluster to use the resources generated by oc-mirror

After you have mirrored your image set to the mirror registry, you must apply the generated **ImageContentSourcePolicy**, **CatalogSource**, and release image signature resources into the cluster.

The **ImageContentSourcePolicy** resource associates the mirror registry with the source registry and redirects image pull requests from the online registries to the mirror registry. The **CatalogSource** resource is used by Operator Lifecycle Manager (OLM) to retrieve information about the available Operators in the mirror registry. The release image signatures are used to verify the mirrored release images.

Prerequisites

- You have mirrored the image set to the registry mirror in the disconnected environment.
- You have access to the cluster as a user with the **cluster-admin** role.

Procedure

1. Log in to the OpenShift CLI as a user with the **cluster-admin** role.
2. Apply the YAML files from the results directory to the cluster by running the following command:

```
$ oc apply -f ./oc-mirror-workspace/results-1639608409/
```

3. If you mirrored release images, apply the release image signatures to the cluster by running the following command:

```
$ oc apply -f ./oc-mirror-workspace/results-1639608409/release-signatures/
```



NOTE

If you are mirroring Operators instead of clusters, you do not need to run **\$ oc apply -f ./oc-mirror-workspace/results-1639608409/release-signatures/**. Running that command will return an error, as there are no release image signatures to apply.

Verification

1. Verify that the **ImageContentSourcePolicy** resources were successfully installed by running the following command:

```
$ oc get imagecontentsourcepolicy
```

2. Verify that the **CatalogSource** resources were successfully installed by running the following command:

```
$ oc get catalogsource -n openshift-marketplace
```

4.4.9. Updating your mirror registry content

You can update your mirror registry content by updating the image set configuration file and mirroring the image set to the mirror registry. The next time that you run the `oc-mirror` plugin, an image set is generated that only contains new and updated images since the previous execution.

While updating the mirror registry, you must take into account the following considerations:

- Images are pruned from the target mirror registry if they are no longer included in the latest image set that was generated and mirrored. Therefore, ensure that you are updating images for the same combination of the following key components so that only a differential image set is created and mirrored:
 - Image set configuration

- Destination registry
- Storage configuration
- The images can be pruned in case of disk to mirror or mirror to mirror workflow.
- The generated image sets must be pushed to the target mirror registry in sequence. You can derive the sequence number from the file name of the generated image set archive file.
- Do not delete or modify the metadata image that is generated by the oc-mirror plugin.
- If you specified a top-level namespace for the mirror registry during the initial image set creation, then you must use this same namespace every time you run the oc-mirror plugin for the same mirror registry.

For more information about the workflow to update the mirror registry content, see the "High level workflow" section.

4.4.9.1. Mirror registry update examples

This section covers the use cases for updating the mirror registry from disk to mirror.

Example ImageSetConfiguration file that was previously used for mirroring:

```
apiVersion: mirror.openshift.io/v1alpha2
kind: ImageSetConfiguration
storageConfig:
  local:
    path: /home/user/metadata
mirror:
  platform:
    channels:
      - name: stable-4.12
        minVersion: 4.12.1
        maxVersion: 4.12.1
  operators:
    - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.14
      packages:
        - name: rhacs-operator
          channels:
            - name: stable
```

Mirroring a specific OpenShift Container Platform version by pruning the existing images

Updated ImageSetConfiguration file:

```
apiVersion: mirror.openshift.io/v1alpha2
kind: ImageSetConfiguration
storageConfig:
  local:
    path: /home/user/metadata
mirror:
  platform:
    channels:
      - name: stable-4.13 1
```

```

operators:
- catalog: registry.redhat.io/redhat/redhat-operator-index:v4.14
  packages:
  - name: rhacs-operator
    channels:
    - name: stable

```

- 1 Replacing by **stable-4.13** prunes all the images of **stable-4.12**.

Updating to the latest version of an Operator by pruning the existing images

Updated ImageSetConfiguration file:

```

apiVersion: mirror.openshift.io/v1alpha2
kind: ImageSetConfiguration
storageConfig:
  local:
    path: /home/user/metadata
mirror:
  platform:
    channels:
    - name: stable-4.12
      minVersion: 4.12.1
      maxVersion: 4.12.1
  operators:
  - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.14
    packages:
    - name: rhacs-operator
      channels:
      - name: stable 1

```

- 1 Using the same channel without specifying a version prunes the existing images and updates with the latest version of images.

Mirroring a new Operator by pruning the existing Operator

Updated ImageSetConfiguration file:

```

apiVersion: mirror.openshift.io/v1alpha2
kind: ImageSetConfiguration
storageConfig:
  local:
    path: /home/user/metadata
mirror:
  platform:
    channels:
    - name: stable-4.12
      minVersion: 4.12.1
      maxVersion: 4.12.1
  operators:
  - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.14
    packages:

```

```
- name: <new_operator_name> 1
  channels:
  - name: stable
```

- 1** Replacing **rhacs-operator** with **new_operator_name** prunes the Red Hat Advanced Cluster Security for Kubernetes Operator.

Pruning all the OpenShift Container Platform images

Updated ImageSetConfiguration file:

```
apiVersion: mirror.openshift.io/v1alpha2
kind: ImageSetConfiguration
storageConfig:
  local:
    path: /home/user/metadata
mirror:
  platform:
    channels:
  operators:
  - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.14
  packages:
```

Additional resources

- [Image set configuration examples](#)
- [Mirroring an image set in a partially disconnected environment](#)
- [Mirroring an image set in a fully disconnected environment](#)
- [Configuring your cluster to use the resources generated by oc-mirror](#)

4.4.10. Performing a dry run

You can use `oc-mirror` to perform a dry run, without actually mirroring any images. This allows you to review the list of images that would be mirrored, as well as any images that would be pruned from the mirror registry. A dry run also allows you to catch any errors with your image set configuration early or use the generated list of images with other tools to carry out the mirroring operation.

Prerequisites

- You have access to the internet to obtain the necessary container images.
- You have installed the OpenShift CLI (**oc**).
- You have installed the **oc-mirror** CLI plugin.
- You have created the image set configuration file.

Procedure

1. Run the **oc mirror** command with the **--dry-run** flag to perform a dry run:

```
$ oc mirror --config=./imageset-config.yaml \ 1
docker://registry.example:5000 \ 2
--dry-run 3
```

- 1 Pass in the image set configuration file that was created. This procedure assumes that it is named **imageset-config.yaml**.
- 2 Specify the mirror registry. Nothing is mirrored to this registry as long as you use the **--dry-run** flag.
- 3 Use the **--dry-run** flag to generate the dry run artifacts and not an actual image set file.

Example output

```
Checking push permissions for registry.example:5000
Creating directory: oc-mirror-workspace/src/publish
Creating directory: oc-mirror-workspace/src/v2
Creating directory: oc-mirror-workspace/src/charts
Creating directory: oc-mirror-workspace/src/release-signatures
No metadata detected, creating new workspace
wrote mirroring manifests to oc-mirror-workspace/operators.1658342351/manifests-redhat-operator-index

...

info: Planning completed in 31.48s
info: Dry run complete
Writing image mapping to oc-mirror-workspace/mapping.txt
```

2. Navigate into the workspace directory that was generated:

```
$ cd oc-mirror-workspace/
```

3. Review the **mapping.txt** file that was generated.
This file contains a list of all images that would be mirrored.
4. Review the **pruning-plan.json** file that was generated.
This file contains a list of all images that would be pruned from the mirror registry when the image set is published.



NOTE

The **pruning-plan.json** file is only generated if your `oc-mirror` command points to your mirror registry and there are images to be pruned.

4.4.11. Including local OCI Operator catalogs

While mirroring OpenShift Container Platform releases, Operator catalogs, and additional images from a registry to a partially disconnected cluster, you can include Operator catalog images from a local file-based catalog on disk. The local catalog must be in the Open Container Initiative (OCI) format.

The local catalog and its contents are mirrored to your target mirror registry based on the filtering information in the image set configuration file.



IMPORTANT

When mirroring local OCI catalogs, any OpenShift Container Platform releases or additional images that you want to mirror along with the local OCI-formatted catalog must be pulled from a registry.

You cannot mirror OCI catalogs along with an `oc-mirror` image set file on disk.

One example use case for using the OCI feature is if you have a CI/CD system building an OCI catalog to a location on disk, and you want to mirror that OCI catalog along with an OpenShift Container Platform release to your mirror registry.



NOTE

If you used the Technology Preview OCI local catalogs feature for the `oc-mirror` plugin for OpenShift Container Platform 4.12, you can no longer use the OCI local catalogs feature of the `oc-mirror` plugin to copy a catalog locally and convert it to OCI format as a first step to mirroring to a fully disconnected cluster.

Prerequisites

- You have access to the internet to obtain the necessary container images.
- You have installed the OpenShift CLI (**oc**).
- You have installed the **oc-mirror** CLI plugin.

Procedure

1. Create the image set configuration file and adjust the settings as necessary.
The following example image set configuration mirrors an OCI catalog on disk along with an OpenShift Container Platform release and a UBI image from **registry.redhat.io**.

```

kind: ImageSetConfiguration
apiVersion: mirror.openshift.io/v1alpha2
storageConfig:
  local:
    path: /home/user/metadata 1
  mirror:
    platform:
      channels:
        - name: stable-4.16 2
          type: ocp
          graph: false
      operators:
        - catalog: oci:///home/user/oc-mirror/my-oci-catalog 3
          targetCatalog: my-namespace/redhat-operator-index 4
          packages:
            - name: aws-load-balancer-operator
        - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.16 5
          packages:
            - name: rhacs-operator
    additionalImages:
      - name: registry.redhat.io/ubi9/ubi:latest 6

```

- 1 Set the back-end location to save the image set metadata to. This location can be a registry or local directory. It is required to specify **storageConfig** values.
- 2 Optionally, include an OpenShift Container Platform release to mirror from **registry.redhat.io**.
- 3 Specify the absolute path to the location of the OCI catalog on disk. The path must start with **oci://** when using the OCI feature.
- 4 Optionally, specify an alternative namespace and name to mirror the catalog as.
- 5 Optionally, specify additional Operator catalogs to pull from a registry.
- 6 Optionally, specify additional images to pull from a registry.

2. Run the **oc mirror** command to mirror the OCI catalog to a target mirror registry:

```
$ oc mirror --config=./imageset-config.yaml \ 1
docker://registry.example:5000 2
```

- 1 Pass in the image set configuration file. This procedure assumes that it is named **imageset-config.yaml**.
- 2 Specify the registry to mirror the content to. The registry must start with **docker://**. If you specify a top-level namespace for the mirror registry, you must also use this same namespace on subsequent executions.

Optionally, you can specify other flags to adjust the behavior of the OCI feature:

--oci-insecure-signature-policy

Do not push signatures to the target mirror registry.

--oci-registries-config

Specify the path to a TOML-formatted **registries.conf** file. You can use this to mirror from a different registry, such as a pre-production location for testing, without having to change the image set configuration file. This flag only affects local OCI catalogs, not any other mirrored content.

Example registries.conf file

```
[[registry]]
location = "registry.redhat.io:5000"
insecure = false
blocked = false
mirror-by-digest-only = true
prefix = ""
[[registry.mirror]]
location = "preprod-registry.example.com"
insecure = false
```

Next steps

- Configure your cluster to use the resources generated by **oc-mirror**.

Additional resources

- [Configuring your cluster to use the resources generated by oc-mirror](#)

4.4.12. Image set configuration parameters

The oc-mirror plugin requires an image set configuration file that defines what images to mirror. The following table lists the available parameters for the **ImageSetConfiguration** resource.

Table 4.1. ImageSetConfiguration parameters

Parameter	Description	Values
apiVersion	The API version for the ImageSetConfiguration content.	String. For example: mirror.openshift.io/v1alpha2 .
archiveSize	The maximum size, in GiB, of each archive file within the image set.	Integer. For example: 4
mirror	The configuration of the image set.	Object
mirror.additionalImages	The additional images configuration of the image set.	Array of objects. For example: <pre> additionalImages: - name: registry.redhat.io/ubi8/ubi:latest </pre>
mirror.additionalImages.name	The tag or digest of the image to mirror.	String. For example: registry.redhat.io/ubi8/ubi:latest
mirror.blockedImages	The full tag, digest, or pattern of images to block from mirroring.	Array of strings. For example: docker.io/library/alpine
mirror.helm	The helm configuration of the image set. Note that the oc-mirror plugin supports only helm charts that do not require user input when rendered.	Object

Parameter	Description	Values
mirror.helm.local	The local helm charts to mirror.	Array of objects. For example: <pre>local: - name: podinfo path: /test/podinfo- 5.0.0.tar.gz</pre>
mirror.helm.local.name	The name of the local helm chart to mirror.	String. For example: podinfo .
mirror.helm.local.path	The path of the local helm chart to mirror.	String. For example: /test/podinfo-5.0.0.tar.gz .
mirror.helm.repositories	The remote helm repositories to mirror from.	Array of objects. For example: <pre>repositories: - name: podinfo url: https://example.github.io/podinfo charts: - name: podinfo version: 5.0.0</pre>
mirror.helm.repositories.name	The name of the helm repository to mirror from.	String. For example: podinfo .
mirror.helm.repositories.url	The URL of the helm repository to mirror from.	String. For example: https://example.github.io/podinfo .
mirror.helm.repositories.charts	The remote helm charts to mirror.	Array of objects.

Parameter	Description	Values
mirror.helm.repositories.charts.name	The name of the helm chart to mirror.	String. For example: podinfo .
mirror.helm.repositories.charts.version	The version of the named helm chart to mirror.	String. For example: 5.0.0 .
mirror.operators	The Operators configuration of the image set.	Array of objects. For example: <pre> operators: - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.16 packages: - name: elasticsearch-operator minVersion: '2.4.0'</pre>
mirror.operators.catalog	The Operator catalog to include in the image set.	String. For example: registry.redhat.io/redhat/redhat-operator-index:v4.16 .
mirror.operators.full	When true , downloads the full catalog, Operator package, or Operator channel.	Boolean. The default value is false .

Parameter	Description	Values
mirror.operators.packages	The Operator packages configuration.	Array of objects. For example: <pre>operators: - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.16 packages: - name: elasticsearch-operator minVersion: '5.2.3-31'</pre>
mirror.operators.packages.name	The Operator package name to include in the image set	String. For example: elasticsearch-operator .
mirror.operators.packages.channels	The Operator package channel configuration.	Object
mirror.operators.packages.channels.name	The Operator channel name, unique within a package, to include in the image set.	String. For example: fast or stable-v4.16 .
mirror.operators.packages.channels.maxVersion	The highest version of the Operator mirror across all channels in which it exists. See the following note for further information.	String. For example: 5.2.3-31
mirror.operators.packages.channels.minBundle	The name of the minimum bundle to include, plus all bundles in the update graph to the channel head. Set this field only if the named bundle has no semantic version metadata.	String. For example: bundleName
mirror.operators.packages.channels.minVersion	The lowest version of the Operator to mirror across all channels in which it exists. See the following note for further information.	String. For example: 5.2.3-31

Parameter	Description	Values
mirror.operators.packages.maxVersion	The highest version of the Operator to mirror across all channels in which it exists. See the following note for further information.	String. For example: 5.2.3-31 .
mirror.operators.packages.minVersion	The lowest version of the Operator to mirror across all channels in which it exists. See the following note for further information.	String. For example: 5.2.3-31 .
mirror.operators.skipDependencies	If true , dependencies of bundles are not included.	Boolean. The default value is false .
mirror.operators.targetCatalog	An alternative name and optional namespace hierarchy to mirror the referenced catalog as.	String. For example: my-namespace/my-operator-catalog
mirror.operators.targetName	An alternative name to mirror the referenced catalog as. The targetName parameter is deprecated. Use the targetCatalog parameter instead.	String. For example: my-operator-catalog
mirror.operators.targetTag	An alternative tag to append to the targetName or targetCatalog .	String. For example: v1
mirror.platform	The platform configuration of the image set.	Object

Parameter	Description	Values
mirror.platform.architectures	The architecture of the platform release payload to mirror.	<p>Array of strings. For example:</p> <pre>architectures: - amd64 - arm64 - multi - ppc64le - s390x</pre> <p>The default value is amd64. The value multi ensures that the mirroring is supported for all available architectures, eliminating the need to specify individual architectures.</p>
mirror.platform.channels	The platform channel configuration of the image set.	<p>Array of objects. For example:</p> <pre>channels: - name: stable-4.10 - name: stable-4.16</pre>
mirror.platform.channels.full	When true , sets the minVersion to the first release in the channel and the maxVersion to the last release in the channel.	Boolean. The default value is false .
mirror.platform.channels.name	The name of the release channel.	String. For example: stable-4.16
mirror.platform.channels.minVersion	The minimum version of the referenced platform to be mirrored.	String. For example: 4.12.6
mirror.platform.channels.maxVersion	The highest version of the referenced platform to be mirrored.	String. For example: 4.16.1

Parameter	Description	Values
mirror.platform.channels.shortestPath	Toggles shortest path mirroring or full range mirroring.	Boolean. The default value is false .
mirror.platform.channels.type	The type of the platform to be mirrored.	String. For example: ocp or okd . The default is ocp .
mirror.platform.graph	Indicates whether the OSUS graph is added to the image set and subsequently published to the mirror.	Boolean. The default value is false .
storageConfig	The back-end configuration of the image set.	Object
storageConfig.local	The local back-end configuration of the image set.	Object
storageConfig.local.path	The path of the directory to contain the image set metadata.	String. For example: ./path/to/dir/ .
storageConfig.registry	The registry back-end configuration of the image set.	Object
storageConfig.registry.imageURL	The back-end registry URI. Can optionally include a namespace reference in the URI.	String. For example: quay.io/myuser/imageset:metadata .
storageConfig.registry.skipTLS	Optionally skip TLS verification of the referenced back-end registry.	Boolean. The default value is false .

**NOTE**

Using the **minVersion** and **maxVersion** properties to filter for a specific Operator version range can result in a multiple channel heads error. The error message states that there are **multiple channel heads**. This is because when the filter is applied, the update graph of the Operator is truncated.

Operator Lifecycle Manager requires that every Operator channel contains versions that form an update graph with exactly one end point, that is, the latest version of the Operator. When the filter range is applied, that graph can turn into two or more separate graphs or a graph that has more than one end point.

To avoid this error, do not filter out the latest version of an Operator. If you still run into the error, depending on the Operator, either the **maxVersion** property must be increased or the **minVersion** property must be decreased. Because every Operator graph can be different, you might need to adjust these values until the error resolves.

4.4.13. Image set configuration examples

The following **ImageSetConfiguration** file examples show the configuration for various mirroring use cases.

Use case: Including the shortest OpenShift Container Platform update path

The following **ImageSetConfiguration** file uses a local storage backend and includes all OpenShift Container Platform versions along the shortest update path from the minimum version of **4.11.37** to the maximum version of **4.12.15**.

Example ImageSetConfiguration file

```
apiVersion: mirror.openshift.io/v1alpha2
kind: ImageSetConfiguration
storageConfig:
  local:
    path: /home/user/metadata
mirror:
  platform:
    channels:
      - name: stable-4.12
        minVersion: 4.11.37
        maxVersion: 4.12.15
        shortestPath: true
```

Use case: Including all versions of OpenShift Container Platform from a minimum to the latest version for multi-architecture releases

The following **ImageSetConfiguration** file uses a registry storage backend and includes all OpenShift Container Platform versions starting at a minimum version of **4.13.4** to the latest version in the channel. On every invocation of `oc-mirror` with this image set configuration, the latest release of the **stable-4.13** channel is evaluated, so running `oc-mirror` at regular intervals ensures that you automatically receive the latest releases of OpenShift Container Platform images.

By setting the value of **platform.architectures** to **multi**, you can ensure that the mirroring is supported for multi-architecture releases.

Example ImageSetConfiguration file

```

apiVersion: mirror.openshift.io/v1alpha2
kind: ImageSetConfiguration
storageConfig:
  registry:
    imageURL: example.com/mirror/oc-mirror-metadata
    skipTLS: false
mirror:
  platform:
    architectures:
      - "multi"
  channels:
    - name: stable-4.13
      minVersion: 4.13.4
      maxVersion: 4.13.6

```

Use case: Including Operator versions from a minimum to the latest

The following **ImageSetConfiguration** file uses a local storage backend and includes only the Red Hat Advanced Cluster Security for Kubernetes Operator, versions starting at 4.0.1 and later in the **stable** channel.



NOTE

When you specify a minimum or maximum version range, you might not receive all Operator versions in that range.

By default, oc-mirror excludes any versions that are skipped or replaced by a newer version in the Operator Lifecycle Manager (OLM) specification. Operator versions that are skipped might be affected by a CVE or contain bugs. Use a newer version instead. For more information on skipped and replaced versions, see [Creating an update graph with OLM](#).

To receive all Operator versions in a specified range, you can set the **mirror.operators.full** field to **true**.

Example ImageSetConfiguration file

```

apiVersion: mirror.openshift.io/v1alpha2
kind: ImageSetConfiguration
storageConfig:
  local:
    path: /home/user/metadata
mirror:
  operators:
    - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.16
  packages:
    - name: rhacs-operator
      channels:
        - name: stable
          minVersion: 4.0.1

```



NOTE

To specify a maximum version instead of the latest, set the **mirror.operators.packages.channels.maxVersion** field.

Use case: Including the Nutanix CSI Operator

The following **ImageSetConfiguration** file uses a local storage backend and includes the Nutanix CSI Operator, the OpenShift Update Service (OSUS) graph image, and an additional Red Hat Universal Base Image (UBI).

Example ImageSetConfiguration file

```
kind: ImageSetConfiguration
apiVersion: mirror.openshift.io/v1alpha2
storageConfig:
  registry:
    imageURL: mylocalregistry/ocp-mirror/openshift4
    skipTLS: false
  mirror:
    platform:
      channels:
        - name: stable-4.11
          type: ocp
      graph: true
    operators:
      - catalog: registry.redhat.io/redhat/certified-operator-index:v4.16
    packages:
      - name: nutanixcsioperator
        channels:
          - name: stable
    additionalImages:
      - name: registry.redhat.io/ubi9/ubi:latest
```

Use case: Including the default Operator channel

The following **ImageSetConfiguration** file includes the **stable-5.7** and **stable** channels for the OpenShift Elasticsearch Operator. Even if only the packages from the **stable-5.7** channel are needed, the **stable** channel must also be included in the **ImageSetConfiguration** file, because it is the default channel for the Operator. You must always include the default channel for the Operator package even if you do not use the bundles in that channel.

TIP

You can find the default channel by running the following command: **oc mirror list operators --catalog=<catalog_name> --package=<package_name>**.

Example ImageSetConfiguration file

```
apiVersion: mirror.openshift.io/v1alpha2
kind: ImageSetConfiguration
storageConfig:
  registry:
    imageURL: example.com/mirror/oc-mirror-metadata
    skipTLS: false
  mirror:
    operators:
      - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.16
    packages:
      - name: elasticsearch-operator
```

```
channels:
- name: stable-5.7
- name: stable
```

Use case: Including an entire catalog (all versions)

The following **ImageSetConfiguration** file sets the **mirror.operators.full** field to **true** to include all versions for an entire Operator catalog.

Example ImageSetConfiguration file

```
apiVersion: mirror.openshift.io/v1alpha2
kind: ImageSetConfiguration
storageConfig:
  registry:
    imageURL: example.com/mirror/oc-mirror-metadata
    skipTLS: false
mirror:
  operators:
    - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.16
      full: true
```

Use case: Including an entire catalog (channel heads only)

The following **ImageSetConfiguration** file includes the channel heads for an entire Operator catalog.

By default, for each Operator in the catalog, oc-mirror includes the latest Operator version (channel head) from the default channel. If you want to mirror all Operator versions, and not just the channel heads, you must set the **mirror.operators.full** field to **true**.

This example also uses the **targetCatalog** field to specify an alternative namespace and name to mirror the catalog as.

Example ImageSetConfiguration file

```
apiVersion: mirror.openshift.io/v1alpha2
kind: ImageSetConfiguration
storageConfig:
  registry:
    imageURL: example.com/mirror/oc-mirror-metadata
    skipTLS: false
mirror:
  operators:
    - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.16
      targetCatalog: my-namespace/my-operator-catalog
```

Use case: Including arbitrary images and helm charts

The following **ImageSetConfiguration** file uses a registry storage backend and includes helm charts and an additional Red Hat Universal Base Image (UBI).

Example ImageSetConfiguration file

```
apiVersion: mirror.openshift.io/v1alpha2
kind: ImageSetConfiguration
archiveSize: 4
storageConfig:
```

```

registry:
  imageURL: example.com/mirror/oc-mirror-metadata
  skipTLS: false
mirror:
  platform:
    architectures:
      - "s390x"
    channels:
      - name: stable-4.16
  operators:
    - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.16
  helm:
    repositories:
      - name: redhat-helm-charts
        url: https://raw.githubusercontent.com/redhat-developer/redhat-helm-charts/master
    charts:
      - name: ibm-mongodb-enterprise-helm
        version: 0.2.0
  additionalImages:
    - name: registry.redhat.io/ubi9/ubi:latest

```

4.4.14. Command reference for oc-mirror

The following tables describe the **oc mirror** subcommands and flags:

Table 4.2. oc mirror subcommands

Subcommand	Description
completion	Generate the autocompletion script for the specified shell.
describe	Output the contents of an image set.
help	Show help about any subcommand.
init	Output an initial image set configuration template.
list	List available platform and Operator content and their version.
version	Output the oc-mirror version.

Table 4.3. oc mirror flags

Flag	Description
-c, --config <string>	Specify the path to an image set configuration file.
--continue-on-error	If any non image-pull related error occurs, continue and attempt to mirror as much as possible.

Flag	Description
--dest-skip-tls	Disable TLS validation for the target registry.
--dest-use-http	Use plain HTTP for the target registry.
--dry-run	Print actions without mirroring images. Generates mapping.txt and pruning-plan.json files.
--from <string>	Specify the path to an image set archive that was generated by an execution of oc-mirror to load into a target registry.
-h, --help	Show the help.
--ignore-history	Ignore past mirrors when downloading images and packing layers. Disables incremental mirroring and might download more data.
--manifests-only	Generate manifests for ImageContentSourcePolicy objects to configure a cluster to use the mirror registry, but do not actually mirror any images. To use this flag, you must pass in an image set archive with the --from flag.
--max-nested-paths <int>	Specify the maximum number of nested paths for destination registries that limit nested paths. The default is 0 .
--max-per-registry <int>	Specify the number of concurrent requests allowed per registry. The default is 6 .
--oci-insecure-signature-policy	Do not push signatures when mirroring local OCI catalogs (with --include-local-oci-catalogs).
--oci-registries-config	Provide a registries configuration file to specify an alternative registry location to copy from when mirroring local OCI catalogs (with --include-local-oci-catalogs).
--skip-cleanup	Skip removal of artifact directories.
--skip-image-pin	Do not replace image tags with digest pins in Operator catalogs.
--skip-metadata-check	Skip metadata when publishing an image set. This is only recommended when the image set was created with --ignore-history .
--skip-missing	If an image is not found, skip it instead of reporting an error and aborting execution. Does not apply to custom images explicitly specified in the image set configuration.
--skip-pruning	Disable automatic pruning of images from the target mirror registry.
--skip-verification	Skip digest verification.

Flag	Description
--source-skip-tls	Disable TLS validation for the source registry.
--source-use-http	Use plain HTTP for the source registry.
-v, --verbose <int>	Specify the number for the log level verbosity. Valid values are 0 - 9 . The default is 0 .

4.4.15. Additional resources

- [About cluster updates in a disconnected environment](#)

4.5. MIRRORING IMAGES FOR A DISCONNECTED INSTALLATION BY USING THE OC-MIRROR PLUGIN V2

You can run your cluster in a restricted network without direct internet connectivity if you install the cluster from a mirrored set of OpenShift Container Platform container images in a private registry. This registry must be running whenever your cluster is running.

Just as you can use the **oc-mirror** OpenShift CLI (**oc**) plugin, you can also use oc-mirror plugin v2 to mirror images to a mirror registry in your fully or partially disconnected environments. To download the required images from the official Red Hat registries, you must run oc-mirror plugin v2 from a system with internet connectivity.



IMPORTANT

oc-mirror plugin v2 is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

4.5.1. Prerequisites

- You must have a container image registry that supports [Docker V2-2](#) in the location that hosts the OpenShift Container Platform cluster, such as Red Hat Quay.



NOTE

If you use Red Hat Quay, use version 3.6 or later with the oc-mirror plugin. See the documentation on [Deploying the Red Hat Quay Operator on OpenShift Container Platform \(Red Hat Quay documentation\)](#). If you need additional assistance selecting and installing a registry, contact your sales representative or Red Hat Support.

- If you do not have an existing solution for a container image registry, OpenShift Container Platform subscribers receive a mirror registry for Red Hat OpenShift. This mirror registry is

included with your subscription and serves as a small-scale container registry. You can use this registry to mirror the necessary container images of OpenShift Container Platform for disconnected installations.

- Every machine in the provisioned clusters must have access to the mirror registry. If the registry is unreachable, tasks like installation, updating, or routine operations such as workload relocation, might fail. Mirror registries must be operated in a highly available manner, ensuring their availability aligns with the production availability of your OpenShift Container Platform clusters.

High level workflow

The following steps outline the high-level workflow on how to mirror images to a mirror registry by using the `oc-mirror` plugin v2:

1. Create an image set configuration file.
2. Mirror the image set to the target mirror registry by using one of the following workflows:
 - Mirror an image set directly to the target mirror registry (mirror to mirror).
 - Mirror an image set to disk (Mirror-to-Disk), transfer the **tar** file to the target environment, then mirror the image set to the target mirror registry (Disk-to-Mirror).
3. Configure your cluster to use the resources generated by the `oc-mirror` plugin v2.
4. Repeat these steps to update your target mirror registry as necessary.

4.5.2. About `oc-mirror` plugin v2

The `oc-mirror` OpenShift CLI (**oc**) plugin is a single tool that mirrors all required OpenShift Container Platform content and other images to your mirror registry.

To use the new Technology Preview version of `oc-mirror`, add the **--v2** flag to the `oc-mirror` plugin v2 command line.

`oc-mirror` plugin v2 has the following features:

- Verifies that the complete image set specified in the image set config is mirrored to the mirrored registry, regardless of whether the images were previously mirrored or not.
- Uses a cache system instead of metadata.
- Maintains minimal archive sizes by incorporating only new images into the archive.
- Generates mirroring archives with content selected by mirroring date.
- Can generate **ImageDigestMirrorSet** (IDMS), **ImageTagMirrorSet** (ITMS), instead of **ImageContentSourcePolicy** (ICSP) for the full image set, rather than just for the incremental changes.
- Saves filter Operator versions by bundle name.
- Does not perform automatic pruning. V2 now has a **Delete** feature, which grants users more control over deleting images.
- Introduces support for **registries.conf**. This change facilitates mirroring to multiple enclaves while using the same cache.

4.5.2.1. oc-mirror plugin v2 compatibility and support

The oc-mirror plugin v2 is supported for OpenShift Container Platform.



NOTE

On **aarch64**, **ppc64le**, and **s390x** architectures the oc-mirror plugin v2 is supported only for OpenShift Container Platform versions 4.14 and later.

Use the latest available version of the oc-mirror plugin v2 regardless of which versions of OpenShift Container Platform you need to mirror.

4.5.3. Preparing your mirror hosts

To use the oc-mirror plugin v2 for image mirroring, you need to install the plugin and create a file with credentials for container images, enabling you to mirror from Red Hat to your mirror.

4.5.3.1. Installing the oc-mirror OpenShift CLI plugin

Install the oc-mirror OpenShift CLI plugin to manage image sets in disconnected environments.

Prerequisites

- You have installed the OpenShift CLI (**oc**). If you are mirroring image sets in a fully disconnected environment, ensure the following:
 - You have installed the oc-mirror plugin on the host that has internet access.
 - The host in the disconnected environment has access to the target mirror registry.
- You have set the **umask** parameter to **0022** on the operating system that uses oc-mirror.
- You have installed the correct binary for the RHEL version that you are using.

Procedure

1. Download the oc-mirror CLI plugin.
 - a. Navigate to the [Downloads](#) page of the [OpenShift Cluster Manager](#).
 - b. Under the **OpenShift disconnected installation tools** section, click **Download** for **OpenShift Client (oc) mirror plugin** and save the file.

2. Extract the archive:

```
$ tar xvzf oc-mirror.tar.gz
```

3. If necessary, update the plugin file to be executable:

```
$ chmod +x oc-mirror
```



NOTE

Do not rename the **oc-mirror** file.

4. Install the oc-mirror CLI plugin by placing the file in your **PATH**, for example, **/usr/local/bin**:

```
$ sudo mv oc-mirror /usr/local/bin/.
```

Verification

- Verify that the plugin for oc-mirror v2 is successfully installed by running the following command:

```
$ oc mirror --v2 --help
```

4.5.3.2. Configuring credentials that allow images to be mirrored

Create a container image registry credentials file that enables you to mirror images from Red Hat to your mirror.



WARNING

Do not use this image registry credentials file as the pull secret when you install a cluster. If you provide this file when you install cluster, all of the machines in the cluster will have write access to your mirror registry.



WARNING

This process requires that you have write access to a container image registry on the mirror registry and adds the credentials to a registry pull secret.

Prerequisites

- You configured a mirror registry to use in your disconnected environment.
- You identified an image repository location on your mirror registry to mirror images into.
- You provisioned a mirror registry account that allows images to be uploaded to that image repository.

Procedure

Complete the following steps on the installation host:

1. Download your **registry.redhat.io** pull secret from [Red Hat OpenShift Cluster Manager](#) .
2. Make a copy of your pull secret in JSON format:

```
$ cat ./pull-secret | jq . > <path>/<pull_secret_file_in_json> 1
```

- 1 Specify the path to the folder to store the pull secret in and a name for the JSON file that you create.

The contents of the file resemble the following example:

```
{
  "auths": {
    "cloud.openshift.com": {
      "auth": "b3BlbnNo...",
      "email": "you@example.com"
    },
    "quay.io": {
      "auth": "b3BlbnNo...",
      "email": "you@example.com"
    },
    "registry.connect.redhat.com": {
      "auth": "NTE3Njg5Nj...",
      "email": "you@example.com"
    },
    "registry.redhat.io": {
      "auth": "NTE3Njg5Nj...",
      "email": "you@example.com"
    }
  }
}
```

3. Save the file as `$XDG_RUNTIME_DIR/containers/auth.json`.
4. Generate the base64-encoded user name and password or token for your mirror registry:

```
$ echo -n '<user_name>:<password>' | base64 -w0 1
BGVtbYk3ZHAqXs=
```

- 1 For `<user_name>` and `<password>`, specify the user name and password that you configured for your registry.

5. Edit the JSON file and add a section that describes your registry to it:

```
"auths": {
  "<mirror_registry>": { 1
    "auth": "<credentials>", 2
    "email": "you@example.com"
  }
},
```

- 1 Specify the registry domain name, and optionally the port, that your mirror registry uses to serve content. For example, **registry.example.com** or **registry.example.com:8443**
- 2 Specify the base64-encoded user name and password for the mirror registry.

The file resembles the following example:

```

{
  "auths": {
    "registry.example.com": {
      "auth": "BGVtbYk3ZHAtqXs=",
      "email": "you@example.com"
    },
    "cloud.openshift.com": {
      "auth": "b3BlbnNo...",
      "email": "you@example.com"
    },
    "quay.io": {
      "auth": "b3BlbnNo...",
      "email": "you@example.com"
    },
    "registry.connect.redhat.com": {
      "auth": "NTE3Njg5Nj...",
      "email": "you@example.com"
    },
    "registry.redhat.io": {
      "auth": "NTE3Njg5Nj...",
      "email": "you@example.com"
    }
  }
}

```

4.5.4. Mirroring an image set to a mirror registry

Mirroring an image set to a mirror registry ensures that the required images are available in a secure and controlled environment, facilitating smoother deployments, updates, and maintenance tasks.

4.5.4.1. Building the image set configuration

The `oc-mirror` plugin v2 uses the image set configuration as the input file to determine the required images for mirroring.

Example for the `ImageSetConfiguration` input file

```

kind: ImageSetConfiguration
apiVersion: mirror.openshift.io/v2alpha1
mirror:
  platform:
    channels:
      - name: stable-4.13
        minVersion: 4.13.10
        maxVersion: 4.13.10
    graph: true
  operators:
    - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.15
      packages:
        - name: aws-load-balancer-operator
        - name: 3scale-operator
        - name: node-observability-operator
  additionalImages:
    - name: registry.redhat.io/ubi8/ubi:latest

```

```
- name:  
registry.redhat.io/ubi9/ubi@sha256:20f695d2a91352d4eaa25107535126727b5945bff38ed36a3e5959  
0f495046f0
```

4.5.4.2. Mirroring an image set in a partially disconnected environment

You can mirror image sets to a registry using the `oc-mirror` plugin v2 in environments with restricted internet access.

Prerequisites

- You have access to the internet and the mirror registry in the environment where you are running the `oc-mirror` plugin v2.

Procedure

- Mirror the images from the specified image set configuration to a specified registry by running the following command:

```
$ oc mirror -c isc.yaml --workspace file://<file_path> docker://<mirror_registry_url> --v2 1
```

- 1** Specify the URL or address of the mirror registry where the images are stored and from which they need to be deleted.

Verification

1. Navigate to the **cluster-resources** directory within the **working-dir** directory that was generated in the **<file_path>** directory.
2. Verify that the YAML files are present for the **ImageDigestMirrorSet**, **ImageTagMirrorSet** and **CatalogSource** resources.

Next steps

- Configure your cluster to use the resources generated by `oc-mirror` plugin v2.

4.5.4.3. Mirroring an image set in a fully disconnected environment

You can mirror image sets in a fully disconnected environment where the OpenShift Container Platform cluster cannot access the public internet.

1. **Mirror to disk:** Prepare an archive containing the image set for mirroring. Internet access is required.
2. **Manual step:** Transfer the archive to the network of the disconnected mirror registry.
3. **Disk to mirror:** To mirror the image set from the archive to the target disconnected registry, run `oc-mirror` plugin v2 from the environment that has access to the mirror registry.

4.5.4.3.1. Mirroring from mirror to disk

You can use the `oc-mirror` plugin v2 to generate an image set and save the content to disk. You can then transfer the generated image set can to the disconnected environment and mirrored to the target registry.

`oc-mirror` plugin v2 retrieves the container images from the source specified in the image set configuration and packs them into a tar archive in a local directory.

Procedure

- Mirror the images from the specified image set configuration to the disk by running the following command:

```
$ oc mirror -c isc.yaml file://<file_path> --v2 1
```

- 1 Add the required file path.

Verification

1. Navigate to the `<file_path>` directory that was generated.
2. Verify that the archive files have been generated.

Next steps

- Configure your cluster to use the resources generated by `oc-mirror` plugin v2.

4.5.4.3.2. Mirroring from disk to mirror

You can use the `oc-mirror` plugin v2 to mirror image sets from a disk to a target mirror registry.

The `oc-mirror` plugin v2 retrieves container images from a local disk and transfers them to the specified mirror registry.

Procedure

- Process the image set file on the disk and mirror the contents to a target mirror registry by running the following command:

```
$ oc mirror -c isc.yaml --from file://<file_path> docker://<mirror_registry_url> --v2 1
```

- 1 Specify the URL or address of the mirror registry where the images are stored and from which they need to be deleted.

Verification

1. Navigate to the `cluster-resources` directory within the `working-dir` directory that was generated in the `<file_path>` directory.
2. Verify that the YAML files are present for the `ImageDigestMirrorSet`, `ImageTagMirrorSet` and `CatalogSource` resources.

Next steps

- Configure your cluster to use the resources generated by oc-mirror plugin v2.

4.5.5. Additional resources

- [Updating a cluster in a disconnected environment using the OpenShift Update Service](#) .

4.5.6. About custom resources generated by v2

With oc-mirror plugin v2, **ImageDigestMirrorSet** (IDMS) and **ImageTagMirrorSet** (ITMS) are generated by default if at least one image is found to which a tag refers. These sets contain mirrors for images referenced by digest or tag in releases, Operator catalogs and additional images.

The **ImageDigestMirrorSet** (IDMS) links the mirror registry to the source registry and forwards image pull requests using digest specifications. The **ImagetagMirrorSet** (ITMS) resource, however, redirects image pull requests by using image tags.

Operator Lifecycle Manager (OLM) uses the **CatalogSource** resource to retrieve information about the available Operators in the mirror registry.

The OSUS service uses the **UpdateService** resource to provide Cincinnati graph to the disconnected environment.

4.5.6.1. Configuring your cluster to use the resources generated by oc-mirror plugin v2

After you have mirrored your image set to the mirror registry, you must apply the generated **ImageDigestMirrorSet** (IDMS), **ImageTagMirrorSet** (ITMS), **CatalogSource**, and **UpdateService** to the cluster.



IMPORTANT

In oc-mirror plugin v2, the IDMS and ITMS files cover the entire image set, unlike the ICSP files in oc-mirror plugin v1. Therefore, the IDMS and ITMS files contain all images of the set even if you only add new images during incremental mirroring.

Prerequisites

- You have access to the cluster as a user with the **cluster-admin** role.

Procedure

- Apply the YAML files from the results directory to the cluster by running the following command:

```
$ oc apply -f <path_to_oc-mirror_workspace>/working-dir/cluster-resources
```

Verification

1. Verify that the **ImageDigestMirrorSet** resources are successfully installed by running the following command:

```
$ oc get imagedigestmirrorset
```


2. Verify that the **ImageTagMirrorSet** resources are successfully installed by running the following command:

```
$ oc get imagetagmirrorset
```

3. Verify that the **CatalogSource** resources are successfully installed by running the following command:

```
$ oc get catalogsource -n openshift-marketplace
```

4.5.7. Deletion of images from your disconnected environment

Before you can use oc-mirror plugin v2, you must delete previously deployed images. oc-mirror plugin v2 no longer performs automatic pruning.


You must create the **DeletedImageSetConfiguration** file to delete image configuration when using oc-mirror plugin v2. This prevents accidentally deleting necessary or deployed images when making changes with **ImageSetConfig.yaml**.

In the following example, **DeletedImageSetConfiguration** removes the following:

- All images of OpenShift Container Platform release 4.13.3.
- The catalog image **redhat-operator-index v4.12**.
- The **aws-load-balancer-operator** v0.0.1 bundle and all its related images.
- The additional images **ubi** and **ubi-minimal** referenced by their corresponding digests.

Example: DeletedImageSetConfig

```
apiVersion: mirror.openshift.io/v2alpha1
kind: DeletedImageSetConfiguration
delete:
  platform:
    channels:
      - name: stable-4.13
        minVersion: 4.13.3
        maxVersion: 4.13.3
    operators:
      - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.12
    packages:
      - name: aws-load-balancer-operator
        minVersion: 0.0.1
        maxVersion: 0.0.1
  additionalImages:
    - name:
registry.redhat.io/ubi8/ubi@sha256:bce7e9f69fb7d4533447232478fd825811c760288f87a35699f9c8f03
0f2c1a6
      - name: registry.redhat.io/ubi8/ubi-
minimal@sha256:8bedbe742f140108897fb3532068e8316900d9814f399d676ac78b46e740e34e
```



IMPORTANT

Consider using the mirror-to-disk and disk-to-mirror workflows to reduce mirroring issues.

In the image delete workflow, oc-mirror plugin v2 deletes only the manifests of the images, which does not reduce the storage occupied in the registry.

To free up storage space from unnecessary images, such as those with deleted manifests, you must enable the garbage collector on your container registry. With the garbage collector enabled, the registry will delete the image blobs that no longer have references to any manifests, thereby reducing the storage previously occupied by the deleted blobs. Enabling the garbage collector differs depending on your container registry.

4.5.7.1. Deleting the images from disconnected environment

To delete images from a disconnected environment using the oc-mirror plugin v2, follow the procedure.

Procedure

1. Create a YAML file that deletes previous images:

```
$ oc mirror delete --config delete-image-set-config.yaml --workspace
file://<previously_mirrored_work_folder> --v2 --generate docker://<remote_registry>
```

Where:

- **<previously_mirrored_work_folder>**: Use the directory where images were previously mirrored or stored during the mirroring process.
 - **<remote_registry>**: Insert the URL or address of the remote container registry from which images will be deleted.
2. Go to the **<previously_mirrored_work_folder>/delete directory** that was created.
 3. Verify that the **delete-images.yaml** file has been generated.
 4. Manually ensure that each image listed in the file is no longer needed by the cluster and can be safely removed from the registry.
 5. After you generate the **delete** YAML file, delete the images from the remote registry:

```
$ oc mirror delete --v2 --delete-yaml-file <previously_mirrored_work_folder>/delete/delete-
images.yaml docker:/ <remote_registry>
```

Where:

- **<previously_mirrored_work_folder>**: Specify your previously mirrored work folder.



IMPORTANT

When using the mirror-to-mirror procedure, images are not cached locally, so you cannot delete images from a local cache.

4.5.8. Verifying your selected images for mirroring

You can use `oc-mirror` plugin v2 to perform a test run (dry run) that does not actually mirror any images. This enables you to review the list of images that would be mirrored. You can also use a dry run to catch any errors with your image set configuration early. When running a dry run on a mirror-to-disk workflow, the `oc-mirror` plugin v2 checks if all the images within the image set are available in its cache. Any missing images are listed in the **missing.txt** file. When a dry run is performed before mirroring, both **missing.txt** and **mapping.txt** files contain the same list of images.

4.5.8.1. Performing dry run for oc-mirror plugin v2

Verify your image set configuration by performing a dry run without mirroring any images. This ensures your setup is correct and prevents unintended changes.

Procedure

- To perform a test run, run the **oc mirror** command and append the **--dry-run** argument to the command:

```
$ oc mirror -c <image_set_config_yaml> --from file://<oc_mirror_workspace_path>
docker://<mirror_registry_url> --dry-run --v2
```

Where:

- **<image_set_config_yaml>**: Use the image set configuration file that you just created.
- **<oc_mirror_workspace_path>**: Insert the address of the workspace path.
- **<mirror_registry_url>**: Insert the URL or address of the remote container registry from which images will be deleted.

Example output

```
$ oc mirror --config /tmp/isc_dryrun.yaml file://<oc_mirror_workspace_path> --dry-run --v2

[INFO] : :warning: --v2 flag identified, flow redirected to the oc-mirror v2 version. This is
Tech Preview, it is still under development and it is not production ready.
[INFO] : :wave: Hello, welcome to oc-mirror
[INFO] : :gear: setting up the environment for you...
[INFO] : :twisted_rightwards_arrows: workflow mode: mirrorToDisk
[INFO] : :sleuth_or_spy: going to discover the necessary images...
[INFO] : :mag: collecting release images...
[INFO] : :mag: collecting operator images...
[INFO] : :mag: collecting additional images...
[WARN] : :warning: 54/54 images necessary for mirroring are not available in the
cache.
[WARN] : List of missing images in : CLID-19/working-dir/dry-run/missing.txt.
please re-run the mirror to disk process
[INFO] : :page_facing_up: list of all images for mirroring in : CLID-19/working-dir/dry-
run/mapping.txt
[INFO] : mirror time : 9.641091076s
[INFO] : :wave: Goodbye, thank you for using oc-mirror
```

Verification

1. Navigate to the workspace directory that was generated:

```
$ cd <oc_mirror_workspace_path>
```

2. Review the **mapping.txt** and **missing.txt** files that were generated. These files contain a list of all images that would be mirrored.

4.5.9. Benefits of enclave support

Enclave support restricts internal access to a specific part of a network. Unlike a demilitarized zone (DMZ) network, which allows inbound and outbound traffic access through firewall boundaries, enclaves do not cross firewall boundaries.



IMPORTANT

Enclave Support is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

The new enclave support functionality is for scenarios where mirroring is needed for multiple enclaves that are secured behind at least one intermediate disconnected network.

Enclave support has the following benefits:

- You can mirror content for multiple enclaves and centralize it in a single internal registry. Because some customers want to run security checks on the mirrored content, with this setup they can run these checks all at once. The content is then vetted before being mirrored to downstream enclaves.
- You can mirror content directly from the centralized internal registry to enclaves without restarting the mirroring process from the internet for each enclave.
- You can minimize data transfer between network stages, so to ensure that a blob or image is transferred only once from one stage to another.

4.5.9.1. Mirroring to an enclave

When you mirror to an enclave, you must first transfer the necessary images from one or more enclaves into the enterprise central registry.

The central registry is situated within a secure network, specifically a disconnected environment, and is not directly linked to the public internet. But the user must execute **oc mirror** in an environment with access to the public internet.

Procedure

1. Before running oc-mirror plugin v2 in the disconnected environment, create a **registries.conf** file. The TOML format of the file is described in this specification:

**NOTE**

It is recommended to store the file under **\$HOME/.config/containers/registries.conf** or **/etc/containers/registries.conf**.

Example registries.conf

```
[[registry]]
location="registry.redhat.io"
[[registry.mirror]]
location="<enterprise-registry.in>"

[[registry]]
location="quay.io"
[[registry.mirror]]
location="<enterprise-registry.in>"
```

2. Generate a mirror archive.
 - a. To collect all the OpenShift Container Platform content into an archive on the disk under **<file_path>/enterprise-content**, run the following command:

```
$ oc mirror --v2 -c isc.yaml file://<file_path>/enterprise-content
```

Example of isc.yaml:

```
apiVersion: mirror.openshift.io/v2alpha1
kind: ImageSetConfiguration
mirror:
  platform:
    architectures:
      - "amd64"
    channels:
      - name: stable-4.15
        minVersion: 4.15.0
        maxVersion: 4.15.3
```

After the archive is generated, it is transferred to the disconnected environment. The transport mechanism is not part of oc-mirror plugin v2. The enterprise network administrators determine the transfer strategy.

In some cases, the transfer is done manually, in that the disk is physically unplugged from one location, and plugged to another computer in the disconnected environment. In other cases, the Secure File Transfer Protocol (SFTP) or other protocols are used.

3. After the transfer of the archive is done, you can execute oc-mirror plugin v2 again in order to mirror the relevant archive contents to the registry (**enterprise_registry.in** in the example) as demonstrated in the following example:

```
$ oc mirror --v2 -c isc.yaml --from file://<disconnected_environment_file_path>/enterprise-content docker://<enterprise_registry.in>/
```

Where:

- **--from** points to the folder containing the archive. It starts with the **file://**.
 - **docker://** is the destination of the mirroring is the final argument. Because it is a docker registry.
 - **-c (--config)** is a mandatory argument. It enables oc-mirror plugin v2 to eventually mirror only sub-parts of the archive to the registry. One archive might contain several OpenShift Container Platform releases, but the disconnected environment or an enclave might mirror only a few.
4. Prepare the **imageSetConfig** YAML file, which describes the content to mirror to the enclave:

Example isc-enclave.yaml

```
apiVersion: mirror.openshift.io/v2alpha1
kind: ImageSetConfiguration
mirror:
  platform:
    architectures:
      - "amd64"
    channels:
      - name: stable-4.15
        minVersion: 4.15.2
        maxVersion: 4.15.2
```

You must run oc-mirror plugin v2 on a machine with access to the disconnected registry. In the previous example, the disconnected environment, **enterprise-registry.in**, is accessible.

5. Update the graph URL
If you are using **graph:true**, oc-mirror plugin v2 attempts to reach the **cincinnati** API endpoint. Because this environment is disconnected, be sure to export the environment variable **UPDATE_URL_OVERRIDE** to refer to the URL for the OpenShift Update Service (OSUS):

```
$ export UPDATE_URL_OVERRIDE=https://<osus.enterprise.in>/graph
```

For more information on setting up OSUS on an OpenShift cluster, see "Updating a cluster in a disconnected environment using the OpenShift Update Service".

6. Generate a mirror archive from the enterprise registry for the enclave.
To prepare an archive for the **enclave1**, the user executes oc-mirror plugin v2 in the enterprise disconnected environment by using the **imageSetConfiguration** specific for that enclave. This ensures that only images needed by that enclave are mirrored:

```
$ oc mirror --v2 -c isc-enclave.yaml
file:///disk-enc1/
```

This action collects all the OpenShift Container Platform content into an archive and generates an archive on disk.

7. After the archive is generated, it will be transferred to the **enclave1** network. The transport mechanism is not the responsibility of oc-mirror plugin v2.
8. Mirror contents to the enclave registry
After the transfer of the archive is done, the user can execute oc-mirror plugin v2 again in order to mirror the relevant archive contents to the registry.

```
$ oc mirror --v2 -c isc-enclave.yaml --from file://local-disk docker://registry.enc1.in
```

The administrators of the OpenShift Container Platform cluster in **enclave1** are now ready to install or upgrade that cluster.

4.5.10. How filtering works in the operator catalog

oc-mirror plugin v2 selects the list of bundles for mirroring by processing the information in **ImageSetConfig**.

When oc-mirror plugin v2 selects bundles for mirroring, it does not infer Group Version Kind (GVK) or bundle dependencies, omitting them from the mirroring set. Instead, it strictly adheres to the user instructions. You must explicitly specify any required dependent packages and their versions.

Bundle versions typically use semantic versioning standards (SemVer), and you can sort bundles within a channel by version. You can select bundles that fall within a specific range in the **ImageSetConfig**.

This selection algorithm ensures consistent outcomes compared to oc-mirror plugin v1. However, it does not include upgrade graph details, such as **replaces**, **skip**, and **skipRange**. This approach differs from the OLM algorithm. It might mirror more bundles than necessary for upgrading a cluster because of potentially shorter upgrade paths between the **minVersion** and **maxVersion**.

Table 4.4. Use the following table to see what bundle versions are included in different scenarios:

ImageSetConfig operator filtering	Expected bundle versions
Scenario 1 <pre>mirror: operators: - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.10</pre>	For each package in the catalog, 1 bundle, corresponding to the head version of the default channel for that package.
Scenario 2 <pre>mirror: operators: - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.10 full: true</pre>	All bundles of all channels of the specified catalog

ImageSetConfig operator filtering	Expected bundle versions
<p>Scenario 3</p> <pre> mirror: operators: - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.10 packages: - name: compliance-operator </pre>	<p>One bundle, corresponding to the head version of the default channel for that package</p>
<p>Scenario 4</p> <pre> mirror: operators: - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.10 full: true packages: - name: elasticsearch-operator </pre>	<p>All bundles of all channels for the packages specified</p>
<p>Scenario 5</p> <pre> mirror: operators: - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.15 packages: - name: compliance-operator minVersion: 5.6.0 </pre>	<p>All bundles in the default channel, from the minVersion, up to the channel head for that package that do not rely on the shortest path from upgrade the graph.</p>

ImageSetConfig operator filtering	Expected bundle versions
<p>Scenario 6</p> <pre> mirror: operators: - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.15 packages: - name: compliance-operator maxVersion: 6.0.0 </pre>	<p>All bundles in the default channel that are lower than the maxVersion for that package.</p>
<p>Scenario 7</p> <pre> mirror: operators: - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.15 packages: - name: compliance-operator minVersion: 5.6.0 maxVersion: 6.0.0 </pre>	<p>All bundles in the default channel, between the minVersion and maxVersion for that package. The head of the channel is not included, even if multiple channels are included in the filtering.</p>
<p>Scenario 8</p> <pre> mirror: operators: - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.15 packages: - name: compliance-operator channels - name: stable </pre>	<p>The head bundle for the selected channel of that package.</p>

ImageSetConfig operator filtering	Expected bundle versions
<p>Scenario 9</p> <pre> mirror: operators: - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.10 full: true - packages: - name: elasticsearch-operator channels: - name: 'stable-v0' </pre>	<p>All bundles for the specified packages and channels.</p>
<p>Scenario 10</p> <pre> mirror: operators: - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.15 packages: - name: compliance-operator channels - name: stable - name: stable-5.5 </pre>	<p>The head bundle for each selected channel of that package.</p>
<p>Scenario 11</p> <pre> mirror: operators: - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.15 packages: - name: compliance-operator channels - name: stable minVersion: 5.6.0 </pre>	<p>Within the selected channel of that package, all versions starting with the minVersion up to the channel head. This scenario does not rely on the shortest path from the upgrade graph.</p>

ImageSetConfig operator filtering	Expected bundle versions
<p>Scenario 12</p> <pre> mirror: operators: - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.15 packages: - name: compliance-operator channels: - name: stable maxVersion: 6.0.0 </pre>	<p>Within the selected channel of that package, all versions up to the maxVersion (not relying on the shortest path from the upgrade graph). The head of the channel is not included, even if multiple channels are included in the filtering.</p>
<p>Scenario 13</p> <pre> mirror: operators: - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.15 packages: - name: compliance-operator channels: - name: stable minVersion: 5.6.0 maxVersion: 6.0.0 </pre>	<p>Within the selected channel of that package, all versions between the minVersion and maxVersion, not relying on the shortest path from the upgrade graph. The head of the channel is not included, even if multiple channels are included in the filtering.</p>
<p>Scenario 14</p> <pre> mirror: operators: - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.14 packages: - name: aws-load-balancer-operator bundles: - name: aws-load-balancer-operator.v1.1.0 - name: 3scale-operator bundles: - name: 3scale-operator.v0.10.0-mas </pre>	<p>Only the bundles specified for each package are included in the filtering.</p>

ImageSetConfig operator filtering	Expected bundle versions
<p>Scenario 15</p> <pre> mirror: operators: - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.15 packages: - name: compliance-operator channels - name: stable minVersion: 5.6.0 maxVersion: 6.0.0 </pre>	<p>Do not use this scenario. filtering by channel and by package with a minVersion or maxVersion is not allowed.</p>
<p>Scenario 16</p> <pre> mirror: operators: - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.15 packages: - name: compliance-operator channels - name: stable minVersion: 5.6.0 maxVersion: 6.0.0 </pre>	<p>Do not use this scenario. You cannot filter using full:true and the minVersion or maxVersion.</p>
<p>Scenario 17</p> <pre> mirror: operators: - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.15 full: true packages: - name: compliance-operator channels - name: stable minVersion: 5.6.0 maxVersion: 6.0.0 </pre>	<p>Do not use this scenario. You cannot filter using full:true and the minVersion or maxVersion.</p>

4.5.11. ImageSet configuration parameters for oc-mirror plugin v2

The `oc-mirror` plugin v2 requires an image set configuration file that defines what images to mirror. The following table lists the available parameters for the **ImageSetConfiguration** resource.



NOTE

Using the **minVersion** and **maxVersion** properties to filter for a specific Operator version range can result in a multiple channel heads error. The error message states that there are **multiple channel heads**. This is because when the filter is applied, the update graph of the Operator is truncated.

OLM requires that every Operator channel contains versions that form an update graph with exactly one end point, that is, the latest version of the Operator. When the filter range is applied, that graph can turn into two or more separate graphs or a graph that has more than one end point.

To avoid this error, do not filter out the latest version of an Operator. If you still run into the error, depending on the Operator, either the **maxVersion** property must be increased or the **minVersion** property must be decreased. Because every Operator graph can be different, you might need to adjust these values until the error resolves.

Table 4.5. **ImageSetConfiguration** parameters

Parameter	Description	Values
apiVersion	The API version of the ImageSetConfiguration content.	String Example: mirror.openshift.io/v2alpha1
archiveSize	The maximum size, in GiB, of each archive file within the image set.	Integer Example: 4
mirror	The configuration of the image set.	Object
mirror.additionalImages	The additional images configuration of the image set.	Array of objects Example: <pre>additionalImages: - name: registry.redhat.io/ubi8/ubi:latest</pre>
mirror.additionalImages.name	The tag or digest of the image to mirror.	String Example: registry.redhat.io/ubi8/ubi:latest
mirror.blockedImages	The full tag, digest, or pattern of images to block from mirroring.	Array of strings Example: docker.io/library/alpine

Parameter	Description	Values
mirror.operators	The Operators configuration of the image set.	Array of objects Example: <pre>operators: - catalog: registry.redhat.io/redhat/redhat-operator-index:4.16 packages: - name: elasticsearch-operator minVersion: '2.4.0'</pre>
mirror.operators.catalog	The Operator catalog to include in the image set.	String Example: registry.redhat.io/redhat/redhat-operator-index:v4.15
mirror.operators.full	When true , downloads the full catalog, Operator package, or Operator channel.	Boolean The default value is false .
mirror.operators.packages	The Operator packages configuration.	Array of objects Example: <pre>operators: - catalog: registry.redhat.io/redhat/redhat-operator-index:4.16 packages: - name: elasticsearch-operator minVersion: '5.2.3-31'</pre>
mirror.operators.packages.name	The Operator package name to include in the image set.	String Example: elasticsearch-operator

Parameter	Description	Values
mirror.operators.packages.channels	Operator package channel configuration	Object
mirror.operators.packages.channels.name	The Operator channel name, unique within a package, to include in the image set.	String Example: fast or stable-v4.15
mirror.operators.packages.channels.maxVersion	The highest version of the Operator mirror across all channels in which it exists.	String Example: 5.2.3-31
mirror.operators.packages.channels.minVersion	The lowest version of the Operator to mirror across all channels in which it exists	String Example: 5.2.3-31
mirror.operators.packages.maxVersion	The highest version of the Operator to mirror across all channels in which it exists.	String Example: 5.2.3-31
mirror.operators.packages.minVersion	The lowest version of the Operator to mirror across all channels in which it exists.	String Example: 5.2.3-31
mirror.operators.packages.bundles	Selected bundles configuration	Array of objects Example: <pre>operators: - catalog: registry.redhat.io/redhat/redhat-operator-index:4.16 packages: - name: 3scale- operator bundles: - name: 3scale- operator.v0.10.0-mas</pre>
mirror.operators.packages.bundles.name	Name of the bundle selected for mirror (as it appears in the catalog).	String Example : 3scale-operator.v0.10.0-mas

Parameter	Description	Values
mirror.operators.targetCatalog	An alternative name and optional namespace hierarchy to mirror the referenced catalog as	String Example: my-namespace/my-operator-catalog
mirror.operators.targetCatalogSourceTemplate	Path on disk for a template to use to complete catalogSource custom resource generated by oc-mirror plugin v2.	String Example: /tmp/catalog-source_template.yaml Example of a template file: <pre>apiVersion: operators.coreos.com/v2alpha1 kind: CatalogSource metadata: name: discarded namespace: openshift-marketplace spec: image: discarded sourceType: grpc updateStrategy: registryPoll: interval: 30m0s</pre>
mirror.operators.targetTag	An alternative tag to append to the targetName or targetCatalog .	String Example: v1
mirror.platform	The platform configuration of the image set.	Object

Parameter	Description	Values
mirror.platform.architectures	The architecture of the platform release payload to mirror.	<p>Array of strings Example:</p> <pre>architectures: - amd64 - arm64 - multi - ppc64le - s390x</pre> <p>The default value is amd64. The value multi ensures that the mirroring is supported for all available architectures, eliminating the need to specify individual architectures</p>
mirror.platform.channels	The platform channel configuration of the image set.	<p>Array of objects Example:</p> <pre>channels: - name: stable-4.12 - name: stable-4.16</pre>
mirror.platform.channels.full	When true , sets the minVersion to the first release in the channel and the maxVersion to the last release in the channel.	Boolean The default value is false
mirror.platform.channels.name	Name of the release channel	String Example: stable-4.15
mirror.platform.channels.minVersion	The minimum version of the referenced platform to be mirrored.	String Example: 4.12.6
mirror.platform.channels.maxVersion	The highest version of the referenced platform to be mirrored.	String Example: 4.15.1

Parameter	Description	Values
mirror.platform.channels.shortestPath	Toggles shortest path mirroring or full range mirroring.	Boolean The default value is false
mirror.platform.channels.type	Type of the platform to be mirrored	String Example: ocp or okd . The default is ocp .
mirror.platform.graph	Indicates whether the OSUS graph is added to the image set and subsequently published to the mirror.	Boolean The default value is false

4.5.11.1. Delete ImageSet Configuration parameters

To use the oc-mirror plugin v2, you must have delete image set configuration file that defines which images to delete from the mirror registry. The following table lists the available parameters for the **DeleteImageSetConfiguration** resource.

Table 4.6. DeleteImageSetConfiguration parameters

Parameter	Description	Values
apiVersion	The API version for the DeleteImageSetConfiguration content.	String Example: mirror.openshift.io/v2alpha1
delete	The configuration of the image set to delete.	Object
delete.additionalImages	The additional images configuration of the delete image set.	Array of objects Example: <pre>additionalImages: - name: registry.redhat.io/ubi8/ubi:latest</pre>
delete.additionalImages.name	The tag or digest of the image to delete.	String Example: registry.redhat.io/ubi8/ubi:latest

Parameter	Description	Values
delete.operators	The Operators configuration of the delete image set.	Array of objects Example: <pre> operators: - catalog: registry.redhat.io/redhat/redhat-operator-index: {product-version} packages: - name: elasticsearch-operator minVersion: '2.4.0'</pre>
delete.operators.catalog	The Operator catalog to include in the delete image set.	String Example: registry.redhat.io/redhat/redhat-operator-index:v4.15
delete.operators.full	When true, deletes the full catalog, Operator package, or Operator channel.	Boolean The default value is false
delete.operators.packages	Operator packages configuration	Array of objects Example: <pre> operators: - catalog: registry.redhat.io/redhat/redhat-operator-index: {product-version} packages: - name: elasticsearch-operator minVersion: '5.2.3-31'</pre>

Parameter	Description	Values
delete.operators.packages.name	The Operator package name to include in the delete image set.	String Example: elasticsearch-operator
delete.operators.packages.channels	Operator package channel configuration	Object
delete.operators.packages.channels.name	The Operator channel name, unique within a package, to include in the delete image set.	String Example: fast or stable-v4.15
delete.operators.packages.channels.maxVersion	The highest version of the Operator to delete within the selected channel.	String Example: 5.2.3-31
delete.operators.packages.channels.minVersion	The lowest version of the Operator to delete within the selection in which it exists.	String Example: 5.2.3-31
delete.operators.packages.maxVersion	The highest version of the Operator to delete across all channels in which it exists.	String Example: 5.2.3-31
delete.operators.packages.minVersion	The lowest version of the Operator to delete across all channels in which it exists.	String Example: 5.2.3-31

Parameter	Description	Values
delete.operators.packages.bundles	The selected bundles configuration	<p>Array of objects</p> <p>You cannot choose both channels and bundles for the same operator.</p> <p>Example:</p> <pre> operators: - catalog: registry.redhat.io/redhat/redhat-operator-index: {product-version} packages: - name: 3scale-operator bundles: - name: 3scale-operator.v0.10.0-mas </pre>
delete.operators.packages.bundles.name	Name of the bundle selected to delete (as it is displayed in the catalog)	String Example : 3scale-operator.v0.10.0-mas
delete.platform	The platform configuration of the image set	Object
delete.platform.architectures	The architecture of the platform release payload to delete.	<p>Array of strings</p> <p>Example:</p> <pre> architectures: - amd64 - arm64 - multi - ppc64le - s390x </pre> <p>The default value is amd64</p>

Parameter	Description	Values
delete.platform.channels	The platform channel configuration of the image set.	Array of objects Example: <pre>channels: - name: stable-4.12 - name: stable-4.16</pre>
delete.platform.channels.full	When true , sets the minVersion to the first release in the channel and the maxVersion to the last release in the channel.	Boolean The default value is false
delete.platform.channels.name	Name of the release channel	String Example: stable-4.15
delete.platform.channels.minVersion	The minimum version of the referenced platform to be deleted.	String Example: 4.12.6
delete.platform.channels.maxVersion	The highest version of the referenced platform to be deleted.	String Example: 4.15.1
delete.platform.channels.shortestPath	Toggles between deleting the shortest path and deleting the full range.	Boolean The default value is false
delete.platform.channels.type	Type of the platform to be deleted	String Example: ocp or okd The default is ocp
delete.platform.graph	Determines whether the OSUS graph is deleted as well on the mirror registry as well.	Boolean The default value is false

4.5.12. Command reference for oc-mirror plugin v2

The following tables describe the **oc mirror** subcommands and flags for oc-mirror plugin v2:

Table 4.7. Subcommands and flags for the oc-mirror plugin v2:

Subcommand	Description
help	Show help about any subcommand
version	Output the oc-mirror version

Subcommand	Description
delete	Deletes images in remote registry and local cache.

Table 4.8. oc mirror flags

Flag	Description
--authfile	Displays the string path of the authentication file. Default is `\${XDG_RUNTIME_DIR}/containers/auth.json` .
-c, --config <string>	Specifies the path to an image set configuration file.
--dest-tls-verify	Requires HTTPS and verifies certificates when accessing the container registry or daemon.
--dry-run	Prints actions without mirroring images
--from <string>	Specifies the path to an image set archive that was generated by executing oc-mirror plugin v2 to load a target registry.
-h, --help	Displays help
--loglevel	Displays string log levels. Supported values include info, debug, trace, error. The default is info .
-p, --port	Determines the HTTP port used by oc-mirror plugin v2 local storage instance. The default is 55000 .
--max-nested-paths <int>	Specifies the maximum number of nested paths for destination registries that limit nested paths. The default is 0 .
--secure-policy	Default value is false . If you set a non-default value, the command enables signature verification, which is the secure policy for signature verification.
--since	Includes all new content since a specified date (format: yyyy-mm-dd). When not provided, new content since previous mirroring is mirrored.
--src-tls-verify	Requires HTTPS and verifies certificates when accessing the container registry or daemon.
--strict-archive	Default value is false . If you set a value, the command generates archives that are strictly less than the archiveSize that was set in the imageSetConfig custom resource (CR). Mirroring exist in error if a file being archived exceeds archiveSize (GB).
-v, --version	Displays the version for oc-mirror plugin v2.

Flag	Description
--workspace	Determines string oc-mirror plugin v2 workspace where resources and internal artifacts are generated.

- [Configuring your cluster to use the resources generated by oc-mirror](#)

CHAPTER 5. INSTALLING ON ALIBABA CLOUD

5.1. INSTALLING A CLUSTER ON ALIBABA CLOUD BY USING THE ASSISTED INSTALLER

Alibaba Cloud provides a broad range of cloud computing and data storage services to online businesses and global enterprises. You can install an OpenShift Container Platform cluster on Alibaba Cloud using the Assisted Installer.



IMPORTANT

Installing Alibaba Cloud with Assisted Installer is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

5.1.1. Process outline for creating a cluster with the Assisted Installer

The main steps of the installation process are as follows:

1. Create the cluster with the Assisted Installer and download the generated image.
2. Convert the image to **QCOW2** format. For more information, see the following section.
3. Upload the image to the Object Storage Service bucket in Alibaba Cloud.
4. Import the image to the Elastic Compute Service in Alibaba Cloud.
5. Provision the Alibaba Cloud resources:
 - a. In the Virtual Private Cloud (VPC) console, set the networking configurations.
 - b. In the Alibaba Cloud DNS console, define the Domain Name System.
 - c. In the Elastic Compute Service (ECS) console, provision the compute instances.
6. Complete host discovery in the Assisted Installer.
7. Complete the network configurations in Alibaba Cloud.
8. Complete the cluster configuration and installation in the Assisted Installer.

Additional resources

- [Installing OpenShift Container Platform with the Assisted Installer](#)

5.1.2. Converting the discovery image to QCOW2 format

Convert the generated ISO to **QCOW2** format before importing it into Alibaba Cloud.

Prerequisites

- You have created a cluster and downloaded the discovery image in the Assisted Installer.
- You have access to a Linux machine that is outside the cluster, such as your desktop machine.

Procedure

1. Open the command line interface on the Linux machine.
2. Verify that the system has virtualization flags enabled by running the following command:

```
$ grep -e lm -e svm -e vmx /proc/cpuinfo
```

3. Install the **qemu-img** package on a RHEL or Fedora machine by running the following command:

```
$ sudo dnf install -y qemu-img
```



NOTE

If your system uses the **APT** package manager, install the package using the name **qemu-utils** instead.

4. Convert the image to **QCOW2** by running the following command:

```
$ qemu-img convert -O qcow2 ${CLUSTER_NAME}.iso ${CLUSTER_NAME}.qcow2
```

CHAPTER 6. INSTALLING ON AWS

6.1. INSTALLATION METHODS

You can install OpenShift Container Platform on Amazon Web Services (AWS) using installer-provisioned or user-provisioned infrastructure. The default installation type uses installer-provisioned infrastructure, where the installation program provisions the underlying infrastructure for the cluster. You can also install OpenShift Container Platform on infrastructure that you provision. If you do not use infrastructure that the installation program provisions, you must manage and maintain the cluster resources yourself. You can also install OpenShift Container Platform on a single node, which is a specialized installation method that is ideal for edge computing environments.

6.1.1. Installing a cluster on installer-provisioned infrastructure

You can install a cluster on AWS infrastructure that is provisioned by the OpenShift Container Platform installation program, by using one of the following methods:

- **Installing a cluster quickly on AWS** You can install OpenShift Container Platform on AWS infrastructure that is provisioned by the OpenShift Container Platform installation program. You can install a cluster quickly by using the default configuration options.
- **Installing a customized cluster on AWS** You can install a customized cluster on AWS infrastructure that the installation program provisions. The installation program allows for some customization to be applied at the installation stage. Many other customization options are available [post-installation](#).
- **Installing a cluster on AWS with network customizations** You can customize your OpenShift Container Platform network configuration during installation, so that your cluster can coexist with your existing IP address allocations and adhere to your network requirements.
- **Installing a cluster on AWS in a restricted network** You can install OpenShift Container Platform on AWS on installer-provisioned infrastructure by using an internal mirror of the installation release content. You can use this method to install a cluster that does not require an active internet connection to obtain the software components.
- **Installing a cluster on an existing Virtual Private Cloud** You can install OpenShift Container Platform on an existing AWS Virtual Private Cloud (VPC). You can use this installation method if you have constraints set by the guidelines of your company, such as limits when creating new accounts or infrastructure.
- **Installing a private cluster on an existing VPC** You can install a private cluster on an existing AWS VPC. You can use this method to deploy OpenShift Container Platform on an internal network that is not visible to the internet.
- **Installing a cluster on AWS into a government or secret region** OpenShift Container Platform can be deployed into AWS regions that are specifically designed for US government agencies at the federal, state, and local level, as well as contractors, educational institutions, and other US customers that must run sensitive workloads in the cloud.

6.1.2. Installing a cluster on user-provisioned infrastructure

You can install a cluster on AWS infrastructure that you provision, by using one of the following methods:

- **Installing a cluster on AWS infrastructure that you provide** You can install OpenShift Container Platform on AWS infrastructure that you provide. You can use the provided

CloudFormation templates to create stacks of AWS resources that represent each of the components required for an OpenShift Container Platform installation.

- [Installing a cluster on AWS in a restricted network with user-provisioned infrastructure](#) You can install OpenShift Container Platform on AWS infrastructure that you provide by using an internal mirror of the installation release content. You can use this method to install a cluster that does not require an active internet connection to obtain the software components. You can also use this installation method to ensure that your clusters only use container images that satisfy your organizational controls on external content. While you can install OpenShift Container Platform by using the mirrored content, your cluster still requires internet access to use the AWS APIs.

6.1.3. Installing a cluster on a single node

Installing OpenShift Container Platform on a single node alleviates some of the requirements for high availability and large scale clusters. However, you must address the [requirements for installing on a single node](#), and the [additional requirements for installing single-node OpenShift on a cloud provider](#). After addressing the requirements for single node installation, use the [Installing a customized cluster on AWS](#) procedure to install the cluster. The [installing single-node OpenShift manually](#) section contains an exemplary `install-config.yaml` file when installing an OpenShift Container Platform cluster on a single node.

6.1.4. Additional resources

- [Installation process](#)

6.2. CONFIGURING AN AWS ACCOUNT

Before you can install OpenShift Container Platform, you must configure an Amazon Web Services (AWS) account.

6.2.1. Configuring Route 53

To install OpenShift Container Platform, the Amazon Web Services (AWS) account you use must have a dedicated public hosted zone in your Route 53 service. This zone must be authoritative for the domain. The Route 53 service provides cluster DNS resolution and name lookup for external connections to the cluster.

Procedure

1. Identify your domain, or subdomain, and registrar. You can transfer an existing domain and registrar or obtain a new one through AWS or another source.



NOTE

If you purchase a new domain through AWS, it takes time for the relevant DNS changes to propagate. For more information about purchasing domains through AWS, see [Registering Domain Names Using Amazon Route 53](#) in the AWS documentation.

2. If you are using an existing domain and registrar, migrate its DNS to AWS. See [Making Amazon Route 53 the DNS Service for an Existing Domain](#) in the AWS documentation.


3. Create a public hosted zone for your domain or subdomain. See [Creating a Public Hosted Zone](#) in the AWS documentation.
Use an appropriate root domain, such as **openshiftcorp.com**, or subdomain, such as **clusters.openshiftcorp.com**.
4. Extract the new authoritative name servers from the hosted zone records. See [Getting the Name Servers for a Public Hosted Zone](#) in the AWS documentation.
5. Update the registrar records for the AWS Route 53 name servers that your domain uses. For example, if you registered your domain to a Route 53 service in a different accounts, see the following topic in the AWS documentation: [Adding or Changing Name Servers or Glue Records](#).
6. If you are using a subdomain, add its delegation records to the parent domain. This gives Amazon Route 53 responsibility for the subdomain. Follow the delegation procedure outlined by the DNS provider of the parent domain. See [Creating a subdomain that uses Amazon Route 53 as the DNS service without migrating the parent domain](#) in the AWS documentation for an example high level procedure.

6.2.1.1. Ingress Operator endpoint configuration for AWS Route 53

If you install in either Amazon Web Services (AWS) GovCloud (US) US-West or US-East region, the Ingress Operator uses **us-gov-west-1** region for Route53 and tagging API clients.

The Ingress Operator uses <https://tagging.us-gov-west-1.amazonaws.com> as the tagging API endpoint if a tagging custom endpoint is configured that includes the string 'us-gov-east-1'.

For more information on AWS GovCloud (US) endpoints, see the [Service Endpoints](#) in the AWS documentation about GovCloud (US).



IMPORTANT

Private, disconnected installations are not supported for AWS GovCloud when you install in the **us-gov-east-1** region.

Example Route 53 configuration

```
platform:
  aws:
    region: us-gov-west-1
    serviceEndpoints:
      - name: ec2
        url: https://ec2.us-gov-west-1.amazonaws.com
      - name: elasticloadbalancing
        url: https://elasticloadbalancing.us-gov-west-1.amazonaws.com
      - name: route53
        url: https://route53.us-gov.amazonaws.com ①
      - name: tagging
        url: https://tagging.us-gov-west-1.amazonaws.com ②
```

① Route 53 defaults to <https://route53.us-gov.amazonaws.com> for both AWS GovCloud (US) regions.


② Only the US-West region has endpoints for tagging. Omit this parameter if your cluster is in another region.

6.2.2. AWS account limits

The OpenShift Container Platform cluster uses a number of Amazon Web Services (AWS) components, and the default [Service Limits](#) affect your ability to install OpenShift Container Platform clusters. If you use certain cluster configurations, deploy your cluster in certain AWS regions, or run multiple clusters from your account, you might need to request additional resources for your AWS account.

The following table summarizes the AWS components whose limits can impact your ability to install and run OpenShift Container Platform clusters.

Component	Number of clusters available by default	Default AWS limit	Description
Instance Limits	Varies	Varies	<p>By default, each cluster creates the following instances:</p> <ul style="list-style-type: none"> • One bootstrap machine, which is removed after installation • Three control plane nodes • Three worker nodes <p>These instance type counts are within a new account's default limit. To deploy more worker nodes, enable autoscaling, deploy large workloads, or use a different instance type, review your account limits to ensure that your cluster can deploy the machines that you need.</p> <p>In most regions, the worker machines use an m6i.large instance and the bootstrap and control plane machines use m6i.xlarge instances. In some regions, including all regions that do not support these instance types, m5.large and m5.xlarge instances are used instead.</p>

Component	Number of clusters available by default	Default AWS limit	Description
Elastic IPs (EIPs)	0 to 1	5 EIPs per account	<p>To provision the cluster in a highly available configuration, the installation program creates a public and private subnet for each availability zone within a region. Each private subnet requires a NAT Gateway, and each NAT gateway requires a separate elastic IP. Review the AWS region map to determine how many availability zones are in each region. To take advantage of the default high availability, install the cluster in a region with at least three availability zones. To install a cluster in a region with more than five availability zones, you must increase the EIP limit.</p> <div style="display: flex; align-items: flex-start; margin-top: 10px;">  <div> <p>IMPORTANT</p> <p>To use the us-east-1 region, you must increase the EIP limit for your account.</p> </div> </div>
Virtual Private Clouds (VPCs)	5	5 VPCs per region	Each cluster creates its own VPC.
Elastic Load Balancing (ELB/NLB)	3	20 per region	By default, each cluster creates internal and external network load balancers for the master API server and a single Classic Load Balancer for the router. Deploying more Kubernetes Service objects with type LoadBalancer will create additional load balancers .
NAT Gateways	5	5 per availability zone	The cluster deploys one NAT gateway in each availability zone.

Component	Number of clusters available by default	Default AWS limit	Description
Elastic Network Interfaces (ENIs)	At least 12	350 per region	The default installation creates 21 ENIs and an ENI for each availability zone in your region. For example, the us-east-1 region contains six availability zones, so a cluster that is deployed in that zone uses 27 ENIs. Review the AWS region map to determine how many availability zones are in each region. Additional ENIs are created for additional machines and ELB load balancers that are created by cluster usage and deployed workloads.
VPC Gateway	20	20 per account	Each cluster creates a single VPC Gateway for S3 access.
S3 buckets	99	100 buckets per account	Because the installation process creates a temporary bucket and the registry component in each cluster creates a bucket, you can create only 99 OpenShift Container Platform clusters per AWS account.
Security Groups	250	2,500 per account	Each cluster creates 10 distinct security groups.

6.2.3. Required AWS permissions for the IAM user



NOTE

Your IAM user must have the permission **tag:GetResources** in the region **us-east-1** to delete the base cluster resources. As part of the AWS API requirement, the OpenShift Container Platform installation program performs various actions in this region.

When you attach the **AdministratorAccess** policy to the IAM user that you create in Amazon Web Services (AWS), you grant that user all of the required permissions. To deploy all components of an OpenShift Container Platform cluster, the IAM user requires the following permissions:

Example 6.1. Required EC2 permissions for installation

- **ec2:AttachNetworkInterface**
- **ec2:AuthorizeSecurityGroupEgress**
- **ec2:AuthorizeSecurityGroupIngress**
- **ec2:CopyImage**
- **ec2>CreateNetworkInterface**


- **ec2:CreateSecurityGroup**
- **ec2:CreateTags**
- **ec2:CreateVolume**
- **ec2>DeleteSecurityGroup**
- **ec2>DeleteSnapshot**
- **ec2>DeleteTags**
- **ec2:DeregisterImage**
- **ec2:DescribeAccountAttributes**
- **ec2:DescribeAddresses**
- **ec2:DescribeAvailabilityZones**
- **ec2:DescribeDhcpOptions**
- **ec2:DescribeImages**
- **ec2:DescribeInstanceAttribute**
- **ec2:DescribeInstanceCreditSpecifications**
- **ec2:DescribeInstances**
- **ec2:DescribeInstanceTypes**
- **ec2:DescribeInternetGateways**
- **ec2:DescribeKeyPairs**
- **ec2:DescribeNatGateways**
- **ec2:DescribeNetworkAcls**
- **ec2:DescribeNetworkInterfaces**
- **ec2:DescribePrefixLists**
- **ec2:DescribePublicIpv4Pools** (only required if **publicIpv4Pool** is specified in **install-config.yaml**)
- **ec2:DescribeRegions**
- **ec2:DescribeRouteTables**
- **ec2:DescribeSecurityGroupRules**
- **ec2:DescribeSecurityGroups**
- **ec2:DescribeSubnets**
- **ec2:DescribeTags**

- **ec2:DescribeVolumes**
- **ec2:DescribeVpcAttribute**
- **ec2:DescribeVpcClassicLink**
- **ec2:DescribeVpcClassicLinkDnsSupport**
- **ec2:DescribeVpcEndpoints**
- **ec2:DescribeVpcs**
- **ec2:DisassociateAddress** (only required if **publicIpv4Pool** is specified in **install-config.yaml**)
- **ec2:GetEbsDefaultKmsKeyId**
- **ec2:ModifyInstanceAttribute**
- **ec2:ModifyNetworkInterfaceAttribute**
- **ec2:RevokeSecurityGroupEgress**
- **ec2:RevokeSecurityGroupIngress**
- **ec2:RunInstances**
- **ec2:TerminateInstances**

Example 6.2. Required permissions for creating network resources during installation

- **ec2:AllocateAddress**
- **ec2:AssociateAddress**
- **ec2:AssociateDhcpOptions**
- **ec2:AssociateRouteTable**
- **ec2:AttachInternetGateway**
- **ec2:CreateDhcpOptions**
- **ec2:CreateInternetGateway**
- **ec2:CreateNatGateway**
- **ec2:CreateRoute**
- **ec2:CreateRouteTable**
- **ec2:CreateSubnet**
- **ec2:CreateVpc**
- **ec2:CreateVpcEndpoint**

- **ec2:ModifySubnetAttribute**
- **ec2:ModifyVpcAttribute**

**NOTE**

If you use an existing Virtual Private Cloud (VPC), your account does not require these permissions for creating network resources.

Example 6.3. Required Elastic Load Balancing permissions (ELB) for installation

- **elasticloadbalancing:AddTags**
- **elasticloadbalancing:ApplySecurityGroupsToLoadBalancer**
- **elasticloadbalancing:AttachLoadBalancerToSubnets**
- **elasticloadbalancing:ConfigureHealthCheck**
- **elasticloadbalancing>CreateListener**
- **elasticloadbalancing>CreateLoadBalancer**
- **elasticloadbalancing>CreateLoadBalancerListeners**
- **elasticloadbalancing>CreateTargetGroup**
- **elasticloadbalancing>DeleteLoadBalancer**
- **elasticloadbalancing:DeregisterInstancesFromLoadBalancer**
- **elasticloadbalancing:DeregisterTargets**
- **elasticloadbalancing:DescribeInstanceHealth**
- **elasticloadbalancing:DescribeListeners**
- **elasticloadbalancing:DescribeLoadBalancerAttributes**
- **elasticloadbalancing:DescribeLoadBalancers**
- **elasticloadbalancing:DescribeTags**
- **elasticloadbalancing:DescribeTargetGroupAttributes**
- **elasticloadbalancing:DescribeTargetHealth**
- **elasticloadbalancing:ModifyLoadBalancerAttributes**
- **elasticloadbalancing:ModifyTargetGroup**
- **elasticloadbalancing:ModifyTargetGroupAttributes**
- **elasticloadbalancing:RegisterInstancesWithLoadBalancer**
- **elasticloadbalancing:RegisterTargets**

- **elasticloadbalancing:SetLoadBalancerPoliciesOfListener**
- **elasticloadbalancing:SetSecurityGroups**




IMPORTANT

OpenShift Container Platform uses both the ELB and ELBv2 API services to provision load balancers. The permission list shows permissions required by both services. A known issue exists in the AWS web console where both services use the same **elasticloadbalancing** action prefix but do not recognize the same actions. You can ignore the warnings about the service not recognizing certain **elasticloadbalancing** actions.

Example 6.4. Required IAM permissions for installation

- **iam:AddRoleToInstanceProfile**
- **iam:CreateInstanceProfile**
- **iam:CreateRole**
- **iam:DeleteInstanceProfile**
- **iam>DeleteRole**
- **iam>DeleteRolePolicy**
- **iam:GetInstanceProfile**
- **iam:GetRole**
- **iam:GetRolePolicy**
- **iam:GetUser**
- **iam:ListInstanceProfilesForRole**
- **iam:ListRoles**
- **iam:ListUsers**
- **iam:PassRole**
- **iam:PutRolePolicy**
- **iam:RemoveRoleFromInstanceProfile**
- **iam:SimulatePrincipalPolicy**
- **iam:TagInstanceProfile**
- **iam:TagRole**

**NOTE**

If you have not created a load balancer in your AWS account, the IAM user also requires the **iam:CreateServiceLinkedRole** permission.

Example 6.5. Required Route 53 permissions for installation

- **route53:ChangeResourceRecordSets**
- **route53:ChangeTagsForResource**
- **route53:CreateHostedZone**
- **route53>DeleteHostedZone**
- **route53:GetChange**
- **route53:GetHostedZone**
- **route53:ListHostedZones**
- **route53:ListHostedZonesByName**
- **route53:ListResourceRecordSets**
- **route53:ListTagsForResource**
- **route53:UpdateHostedZoneComment**

Example 6.6. Required Amazon Simple Storage Service (S3) permissions for installation

- **s3:CreateBucket**
- **s3>DeleteBucket**
- **s3:GetAccelerateConfiguration**
- **s3:GetBucketAcl**
- **s3:GetBucketCors**
- **s3:GetBucketLocation**
- **s3:GetBucketLogging**
- **s3:GetBucketObjectLockConfiguration**
- **s3:GetBucketPolicy**
- **s3:GetBucketRequestPayment**
- **s3:GetBucketTagging**
- **s3:GetBucketVersioning**

- **s3:GetBucketWebsite**
- **s3:GetEncryptionConfiguration**
- **s3:GetLifecycleConfiguration**
- **s3:GetReplicationConfiguration**
- **s3:ListBucket**
- **s3:PutBucketAcl**
- **s3:PutBucketPolicy**
- **s3:PutBucketTagging**
- **s3:PutEncryptionConfiguration**

Example 6.7. S3 permissions that cluster Operators require

- **s3>DeleteObject**
- **s3:GetObject**
- **s3:GetObjectAcl**
- **s3:GetObjectTagging**
- **s3:GetObjectVersion**
- **s3:PutObject**
- **s3:PutObjectAcl**
- **s3:PutObjectTagging**

Example 6.8. Required permissions to delete base cluster resources

- **autoscaling:DescribeAutoScalingGroups**
- **ec2>DeleteNetworkInterface**
- **ec2>DeletePlacementGroup**
- **ec2>DeleteVolume**
- **elasticloadbalancing>DeleteTargetGroup**
- **elasticloadbalancing:DescribeTargetGroups**
- **iam>DeleteAccessKey**
- **iam>DeleteUser**
- **iam>DeleteUserPolicy**

- **iam:ListAttachedRolePolicies**
- **iam:ListInstanceProfiles**
- **iam:ListRolePolicies**
- **iam:ListUserPolicies**
- **s3:DeleteObject**
- **s3:ListBucketVersions**
- **tag:GetResources**

Example 6.9. Required permissions to delete network resources

- **ec2:DeleteDhcpOptions**
- **ec2:DeleteInternetGateway**
- **ec2:DeleteNatGateway**
- **ec2:DeleteRoute**
- **ec2:DeleteRouteTable**
- **ec2:DeleteSubnet**
- **ec2:DeleteVpc**
- **ec2:DeleteVpcEndpoints**
- **ec2:DetachInternetGateway**
- **ec2:DisassociateRouteTable**
- **ec2:ReleaseAddress**
- **ec2:ReplaceRouteTableAssociation**



NOTE

If you use an existing VPC, your account does not require these permissions to delete network resources. Instead, your account only requires the **tag:UntagResources** permission to delete network resources.

Example 6.10. Optional permissions for installing a cluster with a custom Key Management Service (KMS) key

- **kms:CreateGrant**
- **kms:Decrypt**
- **kms:DescribeKey**


- **kms:Encrypt**
- **kms:GenerateDataKey**
- **kms:GenerateDataKeyWithoutPlainText**
- **kms:ListGrants**
- **kms:RevokeGrant**

Example 6.11. Required permissions to delete a cluster with shared instance roles

- **iam:UntagRole**

Example 6.12. Additional IAM and S3 permissions that are required to create manifests

- **iam:GetUserPolicy**
- **iam:ListAccessKeys**
- **iam:PutUserPolicy**
- **iam:TagUser**
- **s3:AbortMultipartUpload**
- **s3:GetBucketPublicAccessBlock**
- **s3:ListBucket**
- **s3:ListBucketMultipartUploads**
- **s3:PutBucketPublicAccessBlock**
- **s3:PutLifecycleConfiguration**



NOTE

If you are managing your cloud provider credentials with mint mode, the IAM user also requires the **iam:CreateAccessKey** and **iam:CreateUser** permissions.

Example 6.13. Optional permissions for instance and quota checks for installation

- **ec2:DescribeInstanceTypeOfferings**
- **servicequotas:ListAWSDefaultServiceQuotas**

Example 6.14. Optional permissions for the cluster owner account when installing a cluster on a shared VPC

- **sts:AssumeRole**

Example 6.15. Required permissions for enabling Bring your own public IPv4 addresses (BYOIP) feature for installation

- **ec2:DescribePublicIpv4Pools**
- **ec2:DisassociateAddress**

6.2.4. Creating an IAM user

Each Amazon Web Services (AWS) account contains a root user account that is based on the email address you used to create the account. This is a highly-privileged account, and it is recommended to use it for only initial account and billing configuration, creating an initial set of users, and securing the account.

Before you install OpenShift Container Platform, create a secondary IAM administrative user. As you complete the [Creating an IAM User in Your AWS Account](#) procedure in the AWS documentation, set the following options:

Procedure

1. Specify the IAM user name and select **Programmatic access**.
2. Attach the **AdministratorAccess** policy to ensure that the account has sufficient permission to create the cluster. This policy provides the cluster with the ability to grant credentials to each OpenShift Container Platform component. The cluster grants the components only the credentials that they require.



NOTE

While it is possible to create a policy that grants the all of the required AWS permissions and attach it to the user, this is not the preferred option. The cluster will not have the ability to grant additional credentials to individual components, so the same credentials are used by all components.

3. Optional: Add metadata to the user by attaching tags.
4. Confirm that the user name that you specified is granted the **AdministratorAccess** policy.
5. Record the access key ID and secret access key values. You must use these values when you configure your local machine to run the installation program.



IMPORTANT

You cannot use a temporary session token that you generated while using a multi-factor authentication device to authenticate to AWS when you deploy a cluster. The cluster continues to use your current AWS credentials to create AWS resources for the entire life of the cluster, so you must use key-based, long-term credentials.

6.2.5. IAM Policies and AWS authentication

By default, the installation program creates instance profiles for the bootstrap, control plane, and compute instances with the necessary permissions for the cluster to operate.

However, you can create your own IAM roles and specify them as part of the installation process. You might need to specify your own roles to deploy the cluster or to manage the cluster after installation. For example:

- Your organization's security policies require that you use a more restrictive set of permissions to install the cluster.
- After the installation, the cluster is configured with an Operator that requires access to additional services.

If you choose to specify your own IAM roles, you can take the following steps:

- Begin with the default policies and adapt as required. For more information, see "Default permissions for IAM instance profiles".
- Use the AWS Identity and Access Management Access Analyzer (IAM Access Analyzer) to create a policy template that is based on the cluster's activity. For more information see, "Using AWS IAM Analyzer to create policy templates".

6.2.5.1. Default permissions for IAM instance profiles

By default, the installation program creates IAM instance profiles for the bootstrap, control plane and worker instances with the necessary permissions for the cluster to operate.

The following lists specify the default permissions for control plane and compute machines:

Example 6.16. Default IAM role permissions for control plane instance profiles

- **ec2:AttachVolume**
- **ec2:AuthorizeSecurityGroupIngress**
- **ec2:CreateSecurityGroup**
- **ec2:CreateTags**
- **ec2:CreateVolume**
- **ec2>DeleteSecurityGroup**
- **ec2>DeleteVolume**
- **ec2:Describe***
- **ec2:DetachVolume**
- **ec2:ModifyInstanceAttribute**
- **ec2:ModifyVolume**
- **ec2:RevokeSecurityGroupIngress**
- **elasticloadbalancing:AddTags**

- **elasticloadbalancing:AttachLoadBalancerToSubnets**
- **elasticloadbalancing:ApplySecurityGroupsToLoadBalancer**
- **elasticloadbalancing:CreateListener**
- **elasticloadbalancing:CreateLoadBalancer**
- **elasticloadbalancing:CreateLoadBalancerPolicy**
- **elasticloadbalancing:CreateLoadBalancerListeners**
- **elasticloadbalancing:CreateTargetGroup**
- **elasticloadbalancing:ConfigureHealthCheck**
- **elasticloadbalancing>DeleteListener**
- **elasticloadbalancing>DeleteLoadBalancer**
- **elasticloadbalancing>DeleteLoadBalancerListeners**
- **elasticloadbalancing>DeleteTargetGroup**
- **elasticloadbalancing:DeregisterInstancesFromLoadBalancer**
- **elasticloadbalancing:DeregisterTargets**
- **elasticloadbalancing:Describe***
- **elasticloadbalancing:DetachLoadBalancerFromSubnets**
- **elasticloadbalancing:ModifyListener**
- **elasticloadbalancing:ModifyLoadBalancerAttributes**
- **elasticloadbalancing:ModifyTargetGroup**
- **elasticloadbalancing:ModifyTargetGroupAttributes**
- **elasticloadbalancing:RegisterInstancesWithLoadBalancer**
- **elasticloadbalancing:RegisterTargets**
- **elasticloadbalancing:SetLoadBalancerPoliciesForBackendServer**
- **elasticloadbalancing:SetLoadBalancerPoliciesOfListener**
- **kms:DescribeKey**

Example 6.17. Default IAM role permissions for compute instance profiles

- **ec2:DescribeInstances**
- **ec2:DescribeRegions**

6.2.5.2. Specifying an existing IAM role

Instead of allowing the installation program to create IAM instance profiles with the default permissions, you can use the **install-config.yaml** file to specify an existing IAM role for control plane and compute instances.

Prerequisites

- You have an existing **install-config.yaml** file.

Procedure

1. Update **compute.platform.aws.iamRole** with an existing role for the compute machines.

Sample install-config.yaml file with an IAM role for compute instances

```
compute:
  - hyperthreading: Enabled
    name: worker
  platform:
    aws:
      iamRole: ExampleRole
```

2. Update **controlPlane.platform.aws.iamRole** with an existing role for the control plane machines.

Sample install-config.yaml file with an IAM role for control plane instances

```
controlPlane:
  hyperthreading: Enabled
  name: master
  platform:
    aws:
      iamRole: ExampleRole
```

3. Save the file and reference it when installing the OpenShift Container Platform cluster.



NOTE

To change or update an IAM account after the cluster has been installed, see [RHOCP 4 AWS cloud-credentials access key is expired](#) (Red Hat Knowledgebase).

Additional resources

- [Deploying the cluster](#)

6.2.5.3. Using AWS IAM Analyzer to create policy templates

The minimal set of permissions that the control plane and compute instance profiles require depends on how the cluster is configured for its daily operation.

One way to determine which permissions the cluster instances require is to use the AWS Identity and Access Management Access Analyzer (IAM Access Analyzer) to create a policy template:

- A policy template contains the permissions the cluster has used over a specified period of time.
- You can then use the template to create policies with fine-grained permissions.

Procedure

The overall process could be:

1. Ensure that CloudTrail is enabled. CloudTrail records all of the actions and events in your AWS account, including the API calls that are required to create a policy template. For more information, see the AWS documentation for [working with CloudTrail](#).
2. Create an instance profile for control plane instances and an instance profile for compute instances. Be sure to assign each role a permissive policy, such as PowerUserAccess. For more information, see the AWS documentation for [creating instance profile roles](#).
3. Install the cluster in a development environment and configure it as required. Be sure to deploy all of applications the cluster will host in a production environment.
4. Test the cluster thoroughly. Testing the cluster ensures that all of the required API calls are logged.
5. Use the IAM Access Analyzer to create a policy template for each instance profile. For more information, see the AWS documentation for [generating policies based on the CloudTrail logs](#).
6. Create and add a fine-grained policy to each instance profile.
7. Remove the permissive policy from each instance profile.
8. Deploy a production cluster using the existing instance profiles with the new policies.



NOTE

You can add [IAM Conditions](#) to your policy to make it more restrictive and compliant with your organization security requirements.

6.2.6. Supported AWS Marketplace regions

Installing an OpenShift Container Platform cluster using an AWS Marketplace image is available to customers who purchase the offer in North America.

While the offer must be purchased in North America, you can deploy the cluster to any of the following supported partitions:

- Public
- GovCloud



NOTE

Deploying a OpenShift Container Platform cluster using an AWS Marketplace image is not supported for the AWS secret regions or China regions.

6.2.7. Supported AWS regions

You can deploy an OpenShift Container Platform cluster to the following regions.

**NOTE**

Your IAM user must have the permission **tag:GetResources** in the region **us-east-1** to delete the base cluster resources. As part of the AWS API requirement, the OpenShift Container Platform installation program performs various actions in this region.

6.2.7.1. AWS public regions

The following AWS public regions are supported:

- **af-south-1** (Cape Town)
- **ap-east-1** (Hong Kong)
- **ap-northeast-1** (Tokyo)
- **ap-northeast-2** (Seoul)
- **ap-northeast-3** (Osaka)
- **ap-south-1** (Mumbai)
- **ap-south-2** (Hyderabad)
- **ap-southeast-1** (Singapore)
- **ap-southeast-2** (Sydney)
- **ap-southeast-3** (Jakarta)
- **ap-southeast-4** (Melbourne)
- **ca-central-1** (Central)
- **ca-west-1** (Calgary)
- **eu-central-1** (Frankfurt)
- **eu-central-2** (Zurich)
- **eu-north-1** (Stockholm)
- **eu-south-1** (Milan)
- **eu-south-2** (Spain)
- **eu-west-1** (Ireland)
- **eu-west-2** (London)
- **eu-west-3** (Paris)
- **il-central-1** (Tel Aviv)
- **me-central-1** (UAE)
- **me-south-1** (Bahrain)

- **sa-east-1** (São Paulo)
- **us-east-1** (N. Virginia)
- **us-east-2** (Ohio)
- **us-west-1** (N. California)
- **us-west-2** (Oregon)

6.2.7.2. AWS GovCloud regions

The following AWS GovCloud regions are supported:

- **us-gov-west-1**
- **us-gov-east-1**

6.2.7.3. AWS SC2S and C2S secret regions

The following AWS secret regions are supported:

- **us-isob-east-1** Secret Commercial Cloud Services (SC2S)
- **us-iso-east-1** Commercial Cloud Services (C2S)

6.2.7.4. AWS China regions

The following AWS China regions are supported:

- **cn-north-1** (Beijing)
- **cn-northwest-1** (Ningxia)

6.2.8. Next steps

- Install an OpenShift Container Platform cluster:
 - [Quickly install a cluster](#) with default options on installer-provisioned infrastructure
 - [Install a cluster with cloud customizations on installer-provisioned infrastructure](#)
 - [Install a cluster with network customizations on installer-provisioned infrastructure](#)
 - [Installing a cluster on user-provisioned infrastructure in AWS by using CloudFormation templates](#)

6.3. INSTALLER-PROVISIONED INFRASTRUCTURE

6.3.1. Preparing to install a cluster on AWS

You prepare to install an OpenShift Container Platform cluster on AWS by completing the following steps:

- Verifying internet connectivity for your cluster.

- [Configuring an AWS account](#).
- Downloading the installation program.

**NOTE**

If you are installing in a disconnected environment, you extract the installation program from the mirrored content. For more information, see [Mirroring images for a disconnected installation](#).

- Installing the OpenShift CLI (**oc**).

**NOTE**

If you are installing in a disconnected environment, install **oc** to the mirror host.

- Generating an SSH key pair. You can use this key pair to authenticate into the OpenShift Container Platform cluster's nodes after it is deployed.
- If the cloud identity and access management (IAM) APIs are not accessible in your environment, or if you do not want to store an administrator-level credential secret in the **kube-system** namespace, [manually creating long-term credentials for AWS](#) or [configuring an AWS cluster to use short-term credentials](#) with Amazon Web Services Security Token Service (AWS STS).

6.3.1.1. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to obtain the images that are necessary to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.

6.3.1.2. Obtaining the installation program

Before you install OpenShift Container Platform, download the installation file on the mirror host.

Prerequisites

- You have a computer that runs Linux or macOS, with 500 MB of local disk space.

Procedure


1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select your infrastructure provider.

3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.



IMPORTANT

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.



IMPORTANT

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.

4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#). This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

6.3.1.3. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.



IMPORTANT

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**. If you are updating a cluster in a disconnected environment, install the **oc** version that you plan to update to.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.

5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**.
To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.

**NOTE**

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your PATH.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:


```
$ oc <command>
```

6.3.1.4. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.

**IMPORTANT**

Do not skip this procedure in production environments, where disaster recovery and debugging is required.

**NOTE**


You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1** Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.

**NOTE**

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.


2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.

**NOTE**

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```

**NOTE**

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

6.3.1.5. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources


- See [About remote health monitoring](#) for more information about the Telemetry service

6.3.2. Installing a cluster on AWS

In OpenShift Container Platform version 4.16, you can install a cluster on Amazon Web Services (AWS) that uses the default configuration options.

6.3.2.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [configured an AWS account](#) to host the cluster.




IMPORTANT

If you have an AWS profile stored on your computer, it must not use a temporary session token that you generated while using a multi-factor authentication device. The cluster continues to use your current AWS credentials to create AWS resources for the entire life of the cluster, so you must use key-based, long-term credentials. To generate appropriate keys, see [Managing Access Keys for IAM Users](#) in the AWS documentation. You can supply the keys when you run the installation program.

- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.

6.3.2.2. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.

Procedure

1. Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

- 1** For **<installation_directory>**, specify the directory name to store the files that the installation program creates.
- 2** To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

When specifying the directory:

- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
- Use an empty directory. Some installation assets, such as bootstrap X.509 certificates, have short expiration intervals, therefore you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

2. Provide values at the prompts:

- a. Optional: Select an SSH key to use to access your cluster machines.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- b. Select **aws** as the platform to target.
- c. If you do not have an Amazon Web Services (AWS) profile stored on your computer, enter the AWS access key ID and secret access key for the user that you configured to run the installation program.

**NOTE**

The AWS access key ID and secret access key are stored in `~/.aws/credentials` in the home directory of the current user on the installation host. You are prompted for the credentials by the installation program if the credentials for the exported profile are not present in the file. Any credentials that you provide to the installation program are stored in the file.

- d. Select the AWS region to deploy the cluster to.
 - e. Select the base domain for the Route 53 service that you configured for your cluster.
 - f. Enter a descriptive name for your cluster.
 - g. Paste the [pull secret from Red Hat OpenShift Cluster Manager](#) .
3. Optional: Remove or disable the **AdministratorAccess** policy from the IAM account that you used to install the cluster.

**NOTE**

The elevated permissions provided by the **AdministratorAccess** policy are required only during installation.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to `<installation_directory>/openshift_install.log`.

**IMPORTANT**

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

Additional resources

- See [Configuration and credential file settings](#) in the AWS documentation for more information about AWS profile and credential configuration.

6.3.2.3. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

6.3.2.4. Logging in to the cluster by using the web console

The **kubeadmin** user exists by default after an OpenShift Container Platform installation. You can log in to your cluster as the **kubeadmin** user by using the OpenShift Container Platform web console.


Prerequisites

- You have access to the installation host.
- You completed a cluster installation and all cluster Operators are available.

Procedure

1. Obtain the password for the **kubeadmin** user from the **kubeadmin-password** file on the installation host:

```
$ cat <installation_directory>/auth/kubeadmin-password
```




NOTE

Alternatively, you can obtain the **kubeadmin** password from the **<installation_directory>/openshift_install.log** log file on the installation host.

2. List the OpenShift Container Platform web console route:

```
$ oc get routes -n openshift-console | grep 'console-openshift'
```



NOTE

Alternatively, you can obtain the OpenShift Container Platform route from the **<installation_directory>/openshift_install.log** log file on the installation host.

Example output

```
console    console-openshift-console.apps.<cluster_name>.<base_domain>    console
https reencrypt/Redirect None
```

3. Navigate to the route detailed in the output of the preceding command in a web browser and log in as the **kubeadmin** user.

Additional resources

- See [Accessing the web console](#) for more details about accessing and understanding the OpenShift Container Platform web console.

6.3.2.5. Next steps

- [Validating an installation.](#)
- [Customize your cluster.](#)
- If necessary, you can [opt out of remote health reporting](#) .
- If necessary, you can [remove cloud provider credentials](#).

6.3.3. Installing a cluster on AWS with customizations

In OpenShift Container Platform version 4.16, you can install a customized cluster on infrastructure that the installation program provisions on Amazon Web Services (AWS). To customize the installation, you modify parameters in the **install-config.yaml** file before you install the cluster.



NOTE

The scope of the OpenShift Container Platform installation configurations is intentionally narrow. It is designed for simplicity and ensured success. You can complete many more OpenShift Container Platform configuration tasks after an installation completes.

6.3.3.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [configured an AWS account](#) to host the cluster.



IMPORTANT

If you have an AWS profile stored on your computer, it must not use a temporary session token that you generated while using a multi-factor authentication device. The cluster continues to use your current AWS credentials to create AWS resources for the entire life of the cluster, so you must use long-term credentials. To generate appropriate keys, see [Managing Access Keys for IAM Users](#) in the AWS documentation. You can supply the keys when you run the installation program.

- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.

6.3.3.2. Obtaining an AWS Marketplace image

If you are deploying an OpenShift Container Platform cluster using an AWS Marketplace image, you must first subscribe through AWS. Subscribing to the offer provides you with the AMI ID that the installation program uses to deploy compute nodes.

Prerequisites

- You have an AWS account to purchase the offer. This account does not have to be the same account that is used to install the cluster.

Procedure

1. Complete the OpenShift Container Platform subscription from the [AWS Marketplace](#).
2. Record the AMI ID for your specific AWS Region. As part of the installation process, you must update the **install-config.yaml** file with this value before deploying the cluster.

Sample **install-config.yaml** file with AWS Marketplace compute nodes

```

apiVersion: v1
baseDomain: example.com
compute:
- hyperthreading: Enabled
  name: worker
  platform:
    aws:
      amiID: ami-06c4d345f7c207239 1
      type: m5.4xlarge
  replicas: 3
metadata:
  name: test-cluster
  platform:
    aws:
      region: us-east-2 2
sshKey: ssh-ed25519 AAAA...
pullSecret: '{"auths": ...}'

```

- 1** The AMI ID from your AWS Marketplace subscription.
- 2** Your AMI ID is associated with a specific AWS Region. When creating the installation configuration file, ensure that you select the same AWS Region that you specified when configuring your subscription.

6.3.3.3. Creating the installation configuration file

You can customize the OpenShift Container Platform cluster you install on Amazon Web Services (AWS).

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster. For a restricted network installation, these files are on your mirror host.
- You have the **imageContentSources** values that were generated during mirror registry creation.
- You have obtained the contents of the certificate for your mirror registry.

Procedure

1. Create the **install-config.yaml** file.
 - a. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install create install-config --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the directory name to store the files that the installation program creates.


When specifying the directory:

- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
 - Use an empty directory. Some installation assets, such as bootstrap X.509 certificates, have short expiration intervals, therefore you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.
- b. At the prompts, provide the configuration details for your cloud:
- i. Optional: Select an SSH key to use to access your cluster machines.

**NOTE**

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- ii. Select **AWS** as the platform to target.
- iii. If you do not have an Amazon Web Services (AWS) profile stored on your computer, enter the AWS access key ID and secret access key for the user that you configured to run the installation program.
- iv. Select the AWS region to deploy the cluster to.
- v. Select the base domain for the Route 53 service that you configured for your cluster.
- vi. Enter a descriptive name for your cluster.

**NOTE**

If you are installing a three-node cluster, be sure to set the **compute.replicas** parameter to **0**. This ensures that the cluster's control planes are schedulable. For more information, see "Installing a three-node cluster on AWS".

2. Edit the **install-config.yaml** file to give the additional information that is required for an installation in a restricted network.
 - a. Update the **pullSecret** value to contain the authentication information for your registry:

```
pullSecret: '{"auths":{"<mirror_host_name>:5000": {"auth": "<credentials>","email": "you@example.com"}}}'
```

For **<mirror_host_name>**, specify the registry domain name that you specified in the certificate for your mirror registry, and for **<credentials>**, specify the base64-encoded user name and password for your mirror registry.

- b. Add the **additionalTrustBundle** parameter and value.

```
additionalTrustBundle: |
  -----BEGIN CERTIFICATE-----
```


Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Bootstrap	RHCOS	4	16 GB	100 GB	300
Control plane	RHCOS	4	16 GB	100 GB	300
Compute	RHCOS, RHEL 8.6 and later [3]	2	8 GB	100 GB	300

1. One vCPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio: $(\text{threads per core} \times \text{cores}) \times \text{sockets} = \text{vCPUs}$.
2. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes which require a 10 ms p99 fsync duration. Note that on many cloud platforms, storage size and IOPS scale together, so you might need to over-allocate storage volume to obtain sufficient performance.
3. As with all user-provisioned installations, if you choose to use RHEL compute machines in your cluster, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Use of RHEL 7 compute machines is deprecated and has been removed in OpenShift Container Platform 4.10 and later.

NOTE

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

6.3.3.3.2. Tested instance types for AWS

The following Amazon Web Services (AWS) instance types have been tested with OpenShift Container Platform.



NOTE

Use the machine types included in the following charts for your AWS instances. If you use an instance type that is not listed in the chart, ensure that the instance size you use matches the minimum resource requirements that are listed in the section named "Minimum resource requirements for cluster installation".

Example 6.18. Machine types based on 64-bit x86 architecture

- **c4.***
- **c5.***
- **c5a.***
- **i3.***
- **m4.***
- **m5.***
- **m5a.***
- **m6a.***
- **m6i.***
- **r4.***
- **r5.***
- **r5a.***
- **r6i.***
- **t3.***
- **t3a.***

6.3.3.3.3. Tested instance types for AWS on 64-bit ARM infrastructures

The following Amazon Web Services (AWS) 64-bit ARM instance types have been tested with OpenShift Container Platform.



NOTE

Use the machine types included in the following charts for your AWS ARM instances. If you use an instance type that is not listed in the chart, ensure that the instance size you use matches the minimum resource requirements that are listed in "Minimum resource requirements for cluster installation".

Example 6.19. Machine types based on 64-bit ARM architecture

- **c6g.***
- **m6g.***

6.3.3.3.4. Sample customized install-config.yaml file for AWS

You can customize the installation configuration file (**install-config.yaml**) to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.

**IMPORTANT**

This sample YAML file is provided for reference only. You must obtain your **install-config.yaml** file by using the installation program and modify it.

```

apiVersion: v1
baseDomain: example.com 1
credentialsMode: Mint 2
controlPlane: 3 4
  hyperthreading: Enabled 5
  name: master
  platform:
    aws:
      zones:
      - us-west-2a
      - us-west-2b
      rootVolume:
        iops: 4000
        size: 500
        type: io1 6
      metadataService:
        authentication: Optional 7
        type: m6i.xlarge
      replicas: 3
compute: 8
- hyperthreading: Enabled 9
  name: worker
  platform:
    aws:
      rootVolume:
        iops: 2000
        size: 500
        type: io1 10
      metadataService:
        authentication: Optional 11
        type: c5.4xlarge
      zones:
      - us-west-2c
      replicas: 3
metadata:
  name: test-cluster 12

```



```

networking:
  clusterNetwork:
    - cidr: 10.128.0.0/14
      hostPrefix: 23
  machineNetwork:
    - cidr: 10.0.0.0/16
  networkType: OVNKubernetes 13
  serviceNetwork:
    - 172.30.0.0/16
platform:
  aws:
    region: us-west-2 14
    propagateUserTags: true 15
    userTags:
      adminContact: jdoe
      costCenter: 7536
    subnets: 16
    - subnet-1
    - subnet-2
    - subnet-3
    amiID: ami-0c5d3e03c0ab9b19a 17
    serviceEndpoints: 18
    - name: ec2
      url: https://vpce-id.ec2.us-west-2.vpce.amazonaws.com
    hostedZone: Z3URY6TWQ91KVV 19
  fips: false 20
  sshKey: ssh-ed25519 AAAA... 21
  pullSecret: '{"auths":{"<local_registry>":{"auth": "<credentials>","email": "you@example.com"}}}' 22
  additionalTrustBundle: | 23
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  imageContentSources: 24
  - mirrors:
    - <local_registry>/<local_repository_name>/release
      source: quay.io/openshift-release-dev/ocp-release
  - mirrors:
    - <local_registry>/<local_repository_name>/release
      source: quay.io/openshift-release-dev/ocp-v4.0-art-dev

```

- 1 12 14 Required. The installation program prompts you for this value.
- 2 Optional: Add this parameter to force the Cloud Credential Operator (CCO) to use the specified mode. By default, the CCO uses the root credentials in the **kube-system** namespace to dynamically try to determine the capabilities of the credentials. For details about CCO modes, see the "About the Cloud Credential Operator" section in the *Authentication and authorization* guide.
- 3 8 15 If you do not provide these parameters and values, the installation program provides the default value.
- 4 The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.

- 5 9** Whether to enable or disable simultaneous multithreading, or **hyperthreading**. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores. You



IMPORTANT

If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance. Use larger instance types, such as **m4.2xlarge** or **m5.2xlarge**, for your machines if you disable simultaneous multithreading.

- 6 10** To configure faster storage for etcd, especially for larger clusters, set the storage type as **io1** and set **iops** to **2000**.

- 7 11** Whether to require the [Amazon EC2 Instance Metadata Service v2](#) (IMDSv2). To require IMDSv2, set the parameter value to **Required**. To allow the use of both IMDSv1 and IMDSv2, set the parameter value to **Optional**. If no value is specified, both IMDSv1 and IMDSv2 are allowed.



NOTE

The IMDS configuration for control plane machines that is set during cluster installation can only be changed by using the AWS CLI. The IMDS configuration for compute machines can be changed by using compute machine sets.

- 13** The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 16** If you provide your own VPC, specify subnets for each availability zone that your cluster uses.
- 17** The ID of the AMI used to boot machines for the cluster. If set, the AMI must belong to the same region as the cluster.
- 18** The AWS service endpoints. Custom endpoints are required when installing to an unknown AWS region. The endpoint URL must use the **https** protocol and the host must trust the certificate.
- 19** The ID of your existing Route 53 private hosted zone. Providing an existing hosted zone requires that you supply your own VPC and the hosted zone is already associated with the VPC prior to installing your cluster. If undefined, the installation program creates a new hosted zone.
- 20** Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.



IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 21 You can optionally provide the **sshKey** value that you use to access the machines in your cluster.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- 22 For **<local_registry>**, specify the registry domain name, and optionally the port, that your mirror registry uses to serve content. For example **registry.example.com** or **registry.example.com:5000**. For **<credentials>**, specify the base64-encoded user name and password for your mirror registry.
- 23 Provide the contents of the certificate file that you used for your mirror registry.
- 24 Provide the **imageContentSources** section from the output of the command to mirror the repository.

6.3.3.3.5. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

- Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
```

```
noProxy: ec2.<aws_region>.amazonaws.com,elasticloadbalancing.
<aws_region>.amazonaws.com,s3.<aws_region>.amazonaws.com 3
additionalTrustBundle: | 4
  -----BEGIN CERTIFICATE-----
  <MY_TRUSTED_CA_CERT>
  -----END CERTIFICATE-----
additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations. If you have added the Amazon **EC2,Elastic Load Balancing**, and **S3** VPC endpoints to your VPC, you must add these endpoints to the **noProxy** field.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.



NOTE

The installation program does not support the proxy **readinessEndpoints** field.



NOTE

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.

**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

6.3.3.4. Alternatives to storing administrator-level secrets in the kube-system project

By default, administrator secrets are stored in the **kube-system** project. If you configured the **credentialsMode** parameter in the **install-config.yaml** file to **Manual**, you must use one of the following alternatives:

- To manage long-term cloud credentials manually, follow the procedure in [Manually creating long-term credentials](#).
- To implement short-term credentials that are managed outside the cluster for individual components, follow the procedures in [Configuring an AWS cluster to use short-term credentials](#).

6.3.3.4.1. Manually creating long-term credentials

The Cloud Credential Operator (CCO) can be put into manual mode prior to installation in environments where the cloud identity and access management (IAM) APIs are not reachable, or the administrator prefers not to store an administrator-level credential secret in the cluster **kube-system** namespace.

Procedure

1. If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```
apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...
```

2. If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

3. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

4. Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
--from=$RELEASE_IMAGE \
```

```
--credentials-requests \  
--included \ 1  
--install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \ 2  
--to=<path_to_directory_for_credentials_requests> 3
```

- 1 The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2 Specify the location of the **install-config.yaml** file.
- 3 Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

This command creates a YAML file for each **CredentialsRequest** object.

Sample **CredentialsRequest** object

```
apiVersion: cloudcredential.openshift.io/v1  
kind: CredentialsRequest  
metadata:  
  name: <component_credentials_request>  
  namespace: openshift-cloud-credential-operator  
  ...  
spec:  
  providerSpec:  
    apiVersion: cloudcredential.openshift.io/v1  
    kind: AWSPProviderSpec  
    statementEntries:  
    - effect: Allow  
      action:  
      - iam:GetUser  
      - iam:GetUserPolicy  
      - iam:ListAccessKeys  
    resource: "*"
  ...
```

5. Create YAML files for secrets in the **openshift-install** manifests directory that you generated previously. The secrets must be stored using the namespace and secret name defined in the **spec.secretRef** for each **CredentialsRequest** object.

Sample **CredentialsRequest** object with secrets

```
apiVersion: cloudcredential.openshift.io/v1  
kind: CredentialsRequest  
metadata:  
  name: <component_credentials_request>  
  namespace: openshift-cloud-credential-operator  
  ...  
spec:  
  providerSpec:  
    apiVersion: cloudcredential.openshift.io/v1  
    kind: AWSPProviderSpec  
    statementEntries:  
    - effect: Allow
```

```

action:
- s3:CreateBucket
- s3>DeleteBucket
resource: "*"
...
secretRef:
name: <component_secret>
namespace: <component_namespace>
...

```

Sample Secret object

```

apiVersion: v1
kind: Secret
metadata:
name: <component_secret>
namespace: <component_namespace>
data:
aws_access_key_id: <base64_encoded_aws_access_key_id>
aws_secret_access_key: <base64_encoded_aws_secret_access_key>

```



IMPORTANT

Before upgrading a cluster that uses manually maintained credentials, you must ensure that the CCO is in an upgradeable state.

6.3.3.4.2. Configuring an AWS cluster to use short-term credentials

To install a cluster that is configured to use the AWS Security Token Service (STS), you must configure the CCO utility and create the required AWS resources for your cluster.

6.3.3.4.2.1. Configuring the Cloud Credential Operator utility

To create and manage cloud credentials from outside of the cluster when the Cloud Credential Operator (CCO) is operating in manual mode, extract and prepare the CCO utility (**ccoctl**) binary.



NOTE

The **ccoctl** utility is a Linux binary that must run in a Linux environment.

Prerequisites

- You have access to an OpenShift Container Platform account with cluster administrator access.
- You have installed the OpenShift CLI (**oc**).
- You have created an AWS account for the **ccoctl** utility to use with the following permissions:

Example 6.20. Required AWS permissions

Required iam permissions

- **iam:CreateOpenIDConnectProvider**
- **iam:CreateRole**

- **iam:DeleteOpenIDConnectProvider**
- **iam>DeleteRole**
- **iam>DeleteRolePolicy**
- **iam:GetOpenIDConnectProvider**
- **iam:GetRole**
- **iam:GetUser**
- **iam:ListOpenIDConnectProviders**
- **iam:ListRolePolicies**
- **iam:ListRoles**
- **iam:PutRolePolicy**
- **iam:TagOpenIDConnectProvider**
- **iam:TagRole**

Required s3 permissions

- **s3:CreateBucket**
- **s3>DeleteBucket**
- **s3>DeleteObject**
- **s3:GetBucketAcl**
- **s3:GetBucketTagging**
- **s3:GetObject**
- **s3:GetObjectAcl**
- **s3:GetObjectTagging**
- **s3:ListBucket**
- **s3:PutBucketAcl**
- **s3:PutBucketPolicy**
- **s3:PutBucketPublicAccessBlock**
- **s3:PutBucketTagging**
- **s3:PutObject**
- **s3:PutObjectAcl**
- **s3:PutObjectTagging**

Required cloudfront permissions

- **cloudfront:ListCloudFrontOriginAccessIdentities**
- **cloudfront:ListDistributions**
- **cloudfront:ListTagsForResource**

If you plan to store the OIDC configuration in a private S3 bucket that is accessed by the IAM identity provider through a public CloudFront distribution URL, the AWS account that runs the **ccoctl** utility requires the following additional permissions:

Example 6.21. Additional permissions for a private S3 bucket with CloudFront

- **cloudfront:CreateCloudFrontOriginAccessIdentity**
- **cloudfront:CreateDistribution**
- **cloudfront>DeleteCloudFrontOriginAccessIdentity**
- **cloudfront>DeleteDistribution**
- **cloudfront:GetCloudFrontOriginAccessIdentity**
- **cloudfront:GetCloudFrontOriginAccessIdentityConfig**
- **cloudfront:GetDistribution**
- **cloudfront:TagResource**
- **cloudfront:UpdateDistribution**

**NOTE**

These additional permissions support the use of the **--create-private-s3-bucket** option when processing credentials requests with the **ccoctl aws create-all** command.


Procedure

1. Set a variable for the OpenShift Container Platform release image by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

2. Obtain the CCO container image from the OpenShift Container Platform release image by running the following command:

```
$ CCO_IMAGE=$(oc adm release info --image-for='cloud-credential-operator'  
$RELEASE_IMAGE -a ~/.pull-secret)
```

**NOTE**

Ensure that the architecture of the **\$RELEASE_IMAGE** matches the architecture of the environment in which you will use the **ccoctl** tool.

3. Extract the **ccoctl** binary from the CCO container image within the OpenShift Container Platform release image by running the following command:

```
$ oc image extract $CCO_IMAGE \
  --file="/usr/bin/ccoctl.<rhel_version>" \
  -a ~/.pull-secret
```

- 1 For **<rhel_version>**, specify the value that corresponds to the version of Red Hat Enterprise Linux (RHEL) that the host uses. If no value is specified, **ccoctl.rhel8** is used by default. The following values are valid:

- **rhel8**: Specify this value for hosts that use RHEL 8.
- **rhel9**: Specify this value for hosts that use RHEL 9.

4. Change the permissions to make **ccoctl** executable by running the following command:

```
$ chmod 775 ccoctl.<rhel_version>
```

Verification

- To verify that **ccoctl** is ready to use, display the help file by running the following command:

```
$ ccoctl --help
```

Output of ccoctl --help

```
OpenShift credentials provisioning tool
```

```
Usage:
```

```
ccoctl [command]
```

```
Available Commands:
```

```
aws      Manage credentials objects for AWS cloud
azure    Manage credentials objects for Azure
gcp      Manage credentials objects for Google cloud
help     Help about any command
ibmcloud Manage credentials objects for IBM Cloud
nutanix  Manage credentials objects for Nutanix
```

```
Flags:
```

```
-h, --help  help for ccoctl
```

```
Use "ccoctl [command] --help" for more information about a command.
```

6.3.3.4.2.2. Creating AWS resources with the Cloud Credential Operator utility

You have the following options when creating AWS resources:

- You can use the **ccoctl aws create-all** command to create the AWS resources automatically. This is the quickest way to create the resources. See [Creating AWS resources with a single command](#).
- If you need to review the JSON files that the **ccoctl** tool creates before modifying AWS resources, or if the process the **ccoctl** tool uses to create AWS resources automatically does not meet the requirements of your organization, you can create the AWS resources individually. See [Creating AWS resources individually](#).

6.3.3.4.2.2.1. Creating AWS resources with a single command

If the process the **ccoctl** tool uses to create AWS resources automatically meets the requirements of your organization, you can use the **ccoctl aws create-all** command to automate the creation of AWS resources.

Otherwise, you can create the AWS resources individually. For more information, see "Creating AWS resources individually".



NOTE

By default, **ccoctl** creates objects in the directory in which the commands are run. To create the objects in a different directory, use the **--output-dir** flag. This procedure uses **<path_to_ccoctl_output_dir>** to refer to this directory.

Prerequisites

You must have:

- Extracted and prepared the **ccoctl** binary.

Procedure


1. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

2. Extract the list of **CredentialsRequest** objects from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included 1 \
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml 2 \
  --to=<path_to_directory_for_credentials_requests> 3
```

- 1** The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2** Specify the location of the **install-config.yaml** file.
- 3** Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

**NOTE**

This command might take a few moments to run.

- Use the **ccoctl** tool to process all **CredentialsRequest** objects by running the following command:

```
$ ccoctl aws create-all \
  --name=<name> \ 1
  --region=<aws_region> \ 2
  --credentials-requests-dir=<path_to_credentials_requests_directory> \ 3
  --output-dir=<path_to_ccoctl_output_dir> \ 4
  --create-private-s3-bucket 5
```

- Specify the name used to tag any cloud resources that are created for tracking.
- Specify the AWS region in which cloud resources will be created.
- Specify the directory containing the files for the component **CredentialsRequest** objects.
- Optional: Specify the directory in which you want the **ccoctl** utility to create objects. By default, the utility creates objects in the directory in which the commands are run.
- Optional: By default, the **ccoctl** utility stores the OpenID Connect (OIDC) configuration files in a public S3 bucket and uses the S3 URL as the public OIDC endpoint. To store the OIDC configuration in a private S3 bucket that is accessed by the IAM identity provider through a public CloudFront distribution URL instead, use the **--create-private-s3-bucket** parameter.

**NOTE**

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

Verification

- To verify that the OpenShift Container Platform secrets are created, list the files in the **<path_to_ccoctl_output_dir>/manifests** directory:

```
$ ls <path_to_ccoctl_output_dir>/manifests
```

Example output


```
cluster-authentication-02-config.yaml
openshift-cloud-credential-operator-cloud-credential-operator-iam-ro-creds-credentials.yaml
openshift-cloud-network-config-controller-cloud-credentials-credentials.yaml
openshift-cluster-api-capamanager-bootstrap-credentials-credentials.yaml
openshift-cluster-csi-drivers-efs-cloud-credentials-credentials.yaml
openshift-image-registry-installer-cloud-credentials-credentials.yaml
openshift-ingress-operator-cloud-credentials-credentials.yaml
openshift-machine-api-aws-cloud-credentials-credentials.yaml
```

You can verify that the IAM roles are created by querying AWS. For more information, refer to AWS documentation on listing IAM roles.

6.3.3.4.2.2.2. Creating AWS resources individually

You can use the **ccoctl** tool to create AWS resources individually. This option might be useful for an organization that shares the responsibility for creating these resources among different users or departments.

Otherwise, you can use the **ccoctl aws create-all** command to create the AWS resources automatically. For more information, see "Creating AWS resources with a single command".



NOTE

By default, **ccoctl** creates objects in the directory in which the commands are run. To create the objects in a different directory, use the **--output-dir** flag. This procedure uses **<path_to_ccoctl_output_dir>** to refer to this directory.

Some **ccoctl** commands make AWS API calls to create or modify AWS resources. You can use the **--dry-run** flag to avoid making API calls. Using this flag creates JSON files on the local file system instead. You can review and modify the JSON files and then apply them with the AWS CLI tool using the **--cli-input-json** parameters.

Prerequisites

- Extract and prepare the **ccoctl** binary.

Procedure

1. Generate the public and private RSA key files that are used to set up the OpenID Connect provider for the cluster by running the following command:

```
$ ccoctl aws create-key-pair
```

Example output

```
2021/04/13 11:01:02 Generating RSA keypair
2021/04/13 11:01:03 Writing private key to /<path_to_ccoctl_output_dir>/serviceaccount-
signer.private
2021/04/13 11:01:03 Writing public key to /<path_to_ccoctl_output_dir>/serviceaccount-
signer.public
2021/04/13 11:01:03 Copying signing key for use by installer
```

where **serviceaccount-signer.private** and **serviceaccount-signer.public** are the generated key files.

This command also creates a private key that the cluster requires during installation in **<path_to_ccoctl_output_dir>/tls/bound-service-account-signing-key.key**.

2. Create an OpenID Connect identity provider and S3 bucket on AWS by running the following command:

```
$ ccoctl aws create-identity-provider \
  --name=<name> 1
```

```
--region=<aws_region> \ 2
--public-key-file=<path_to_ccoctl_output_dir>/serviceaccount-signer.public 3
```

- 1** **<name>** is the name used to tag any cloud resources that are created for tracking.
- 2** **<aws-region>** is the AWS region in which cloud resources will be created.
- 3** **<path_to_ccoctl_output_dir>** is the path to the public key file that the **ccoctl aws create-key-pair** command generated.

Example output

```
2021/04/13 11:16:09 Bucket <name>-oidc created
2021/04/13 11:16:10 OpenID Connect discovery document in the S3 bucket <name>-oidc at
.well-known/openid-configuration updated
2021/04/13 11:16:10 Reading public key
2021/04/13 11:16:10 JSON web key set (JWKS) in the S3 bucket <name>-oidc at keys.json
updated
2021/04/13 11:16:18 Identity Provider created with ARN: arn:aws:iam::
<aws_account_id>:oidc-provider/<name>-oidc.s3.<aws_region>.amazonaws.com
```

where **openid-configuration** is a discovery document and **keys.json** is a JSON web key set file.

This command also creates a YAML configuration file in `/<path_to_ccoctl_output_dir>/manifests/cluster-authentication-02-config.yaml`. This file sets the issuer URL field for the service account tokens that the cluster generates, so that the AWS IAM identity provider trusts the tokens.

3. Create IAM roles for each component in the cluster:
 - a. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

- b. Extract the list of **CredentialsRequest** objects from the OpenShift Container Platform release image:

```
$ oc adm release extract \
--from=$RELEASE_IMAGE \
--credentials-requests \
--included \ 1
--install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \
2
--to=<path_to_directory_for_credentials_requests> 3
```

- 1** The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2** Specify the location of the **install-config.yaml** file.
- 3** Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

- c. Use the **ccoctl** tool to process all **CredentialsRequest** objects by running the following command:

```
$ ccoctl aws create-iam-roles \
  --name=<name> \
  --region=<aws_region> \
  --credentials-requests-dir=<path_to_credentials_requests_directory> \
  --identity-provider-arn=arn:aws:iam::<aws_account_id>:oidc-provider/<name>-oidc.s3.
  <aws_region>.amazonaws.com
```



NOTE

For AWS environments that use alternative IAM API endpoints, such as GovCloud, you must also specify your region with the **--region** parameter.

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

For each **CredentialsRequest** object, **ccoctl** creates an IAM role with a trust policy that is tied to the specified OIDC identity provider, and a permissions policy as defined in each **CredentialsRequest** object from the OpenShift Container Platform release image.

Verification

- To verify that the OpenShift Container Platform secrets are created, list the files in the **<path_to_ccoctl_output_dir>/manifests** directory:

```
$ ls <path_to_ccoctl_output_dir>/manifests
```

Example output

```
cluster-authentication-02-config.yaml
openshift-cloud-credential-operator-cloud-credential-operator-iam-ro-creds-credentials.yaml
openshift-cloud-network-config-controller-cloud-credentials-credentials.yaml
openshift-cluster-api-capamanager-bootstrap-credentials-credentials.yaml
openshift-cluster-csi-drivers-ebs-cloud-credentials-credentials.yaml
openshift-image-registry-installer-cloud-credentials-credentials.yaml
openshift-ingress-operator-cloud-credentials-credentials.yaml
openshift-machine-api-aws-cloud-credentials-credentials.yaml
```

You can verify that the IAM roles are created by querying AWS. For more information, refer to AWS documentation on listing IAM roles.

6.3.3.4.2.3. Incorporating the Cloud Credential Operator utility manifests

To implement short-term security credentials managed outside the cluster for individual components, you must move the manifest files that the Cloud Credential Operator utility (**ccoctl**) created to the correct directories for the installation program.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.

- You have configured the Cloud Credential Operator utility (**ccoctl**).
- You have created the cloud provider resources that are required for your cluster with the **ccoctl** utility.

Procedure

1. If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```
apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...
```

2. If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

3. Copy the manifests that the **ccoctl** utility generated to the **manifests** directory that the installation program created by running the following command:


```
$ cp /<path_to_ccoctl_output_dir>/manifests/* ./manifests/
```

4. Copy the private key that the **ccoctl** utility generated in the **tls** directory to the installation directory by running the following command:

```
$ cp -a /<path_to_ccoctl_output_dir>/tls .
```

6.3.3.5. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.

Procedure

1. Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

- 1 For **<installation_directory>**, specify the location of your customized **./install-config.yaml** file.
- 2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

2. Optional: Remove or disable the **AdministratorAccess** policy from the IAM account that you used to install the cluster.



NOTE

The elevated permissions provided by the **AdministratorAccess** policy are required only during installation.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to **<installation_directory>/openshift_install.log**.



IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

6.3.3.6. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

6.3.3.7. Logging in to the cluster by using the web console

The **kubeadmin** user exists by default after an OpenShift Container Platform installation. You can log in to your cluster as the **kubeadmin** user by using the OpenShift Container Platform web console.


Prerequisites

- You have access to the installation host.
- You completed a cluster installation and all cluster Operators are available.

Procedure

1. Obtain the password for the **kubeadmin** user from the **kubeadmin-password** file on the installation host:

```
$ cat <installation_directory>/auth/kubeadmin-password
```



NOTE

Alternatively, you can obtain the **kubeadmin** password from the **<installation_directory>/openshift_install.log** log file on the installation host.

2. List the OpenShift Container Platform web console route:

```
$ oc get routes -n openshift-console | grep 'console-openshift'
```



NOTE

Alternatively, you can obtain the OpenShift Container Platform route from the **<installation_directory>/openshift_install.log** log file on the installation host.

Example output

```
console    console-openshift-console.apps.<cluster_name>.<base_domain>    console
https reencrypt/Redirect None
```

3. Navigate to the route detailed in the output of the preceding command in a web browser and log in as the **kubeadmin** user.

Additional resources

- See [Accessing the web console](#) for more details about accessing and understanding the OpenShift Container Platform web console.

6.3.3.8. Next steps

- [Validating an installation.](#)
- [Customize your cluster.](#)
- If necessary, you can [opt out of remote health reporting](#) .
- If necessary, you can [remove cloud provider credentials](#) .

6.3.4. Installing a cluster on AWS with network customizations

In OpenShift Container Platform version 4.16, you can install a cluster on Amazon Web Services (AWS) with customized network configuration options. By customizing your network configuration, your cluster

can coexist with existing IP address allocations in your environment and integrate with existing MTU and VXLAN configurations.

You must set most of the network configuration parameters during installation, and you can modify only **kubeProxy** configuration parameters in a running cluster.

6.3.4.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [configured an AWS account](#) to host the cluster.



IMPORTANT

If you have an AWS profile stored on your computer, it must not use a temporary session token that you generated while using a multi-factor authentication device. The cluster continues to use your current AWS credentials to create AWS resources for the entire life of the cluster, so you must use key-based, long-term credentials. To generate appropriate keys, see [Managing Access Keys for IAM Users](#) in the AWS documentation. You can supply the keys when you run the installation program.

- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.

6.3.4.2. Network configuration phases

There are two phases prior to OpenShift Container Platform installation where you can customize the network configuration.

Phase 1

You can customize the following network-related fields in the **install-config.yaml** file before you create the manifest files:

- **networking.networkType**
- **networking.clusterNetwork**
- **networking.serviceNetwork**
- **networking.machineNetwork**

For more information on these fields, refer to *Installation configuration parameters*.



NOTE

Set the **networking.machineNetwork** to match the CIDR that the preferred NIC resides in.



IMPORTANT

The CIDR range **172.17.0.0/16** is reserved by libVirt. You cannot use this range or any range that overlaps with this range for any networks in your cluster.

Phase 2

After creating the manifest files by running **openshift-install create manifests**, you can define a customized Cluster Network Operator manifest with only the fields you want to modify. You can use the manifest to specify advanced network configuration.

You cannot override the values specified in phase 1 in the **install-config.yaml** file during phase 2. However, you can further customize the network plugin during phase 2.

6.3.4.3. Creating the installation configuration file

You can customize the OpenShift Container Platform cluster you install on Amazon Web Services (AWS).

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster. For a restricted network installation, these files are on your mirror host.
- You have the **imageContentSources** values that were generated during mirror registry creation.
- You have obtained the contents of the certificate for your mirror registry.

Procedure

1. Create the **install-config.yaml** file.
 - a. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install create install-config --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the directory name to store the files that the installation program creates.

When specifying the directory:

- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
- Use an empty directory. Some installation assets, such as bootstrap X.509 certificates, have short expiration intervals, therefore you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

- b. At the prompts, provide the configuration details for your cloud:

- i. Optional: Select an SSH key to use to access your cluster machines.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- ii. Select **AWS** as the platform to target.
 - iii. If you do not have an Amazon Web Services (AWS) profile stored on your computer, enter the AWS access key ID and secret access key for the user that you configured to run the installation program.
 - iv. Select the AWS region to deploy the cluster to.
 - v. Select the base domain for the Route 53 service that you configured for your cluster.
 - vi. Enter a descriptive name for your cluster.
2. Edit the **install-config.yaml** file to give the additional information that is required for an installation in a restricted network.

- a. Update the **pullSecret** value to contain the authentication information for your registry:

```
pullSecret: '{"auths":{"<mirror_host_name>:5000": {"auth": "<credentials>","email": "you@example.com"}}}'
```

For **<mirror_host_name>**, specify the registry domain name that you specified in the certificate for your mirror registry, and for **<credentials>**, specify the base64-encoded user name and password for your mirror registry.

- b. Add the **additionalTrustBundle** parameter and value.

```
additionalTrustBundle: |
  -----BEGIN CERTIFICATE-----
  /-----END CERTIFICATE-----
```

The value must be the contents of the certificate file that you used for your mirror registry. The certificate file can be an existing, trusted certificate authority, or the self-signed certificate that you generated for the mirror registry.

- c. Define the subnets for the VPC to install the cluster in:

```
subnets:
- subnet-1
- subnet-2
- subnet-3
```

- d. Add the image content resources, which resemble the following YAML excerpt:

```
imageContentSources:
- mirrors:
```

```

- <mirror_host_name>:5000/<repo_name>/release
  source: quay.io/openshift-release-dev/ocp-release
- mirrors:
- <mirror_host_name>:5000/<repo_name>/release
  source: registry.redhat.io/ocp/release

```

For these values, use the **imageContentSources** that you recorded during mirror registry creation.

- e. Optional: Set the publishing strategy to **Internal**:

```
publish: Internal
```

By setting this option, you create an internal Ingress Controller and a private load balancer.

3. Make any other modifications to the **install-config.yaml** file that you require. For more information about the parameters, see "Installation configuration parameters".
4. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

Additional resources

- [Installation configuration parameters for AWS](#)

6.3.4.3.1. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 6.2. Minimum resource requirements

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Bootstrap	RHCOS	4	16 GB	100 GB	300
Control plane	RHCOS	4	16 GB	100 GB	300
Compute	RHCOS, RHEL 8.6 and later [3]	2	8 GB	100 GB	300

1. One vCPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio: (threads per core × cores) × sockets = vCPUs.
2. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes which require a 10 ms

p99 fsync duration. Note that on many cloud platforms, storage size and IOPS scale together, so you might need to over-allocate storage volume to obtain sufficient performance.

3. As with all user-provisioned installations, if you choose to use RHEL compute machines in your cluster, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Use of RHEL 7 compute machines is deprecated and has been removed in OpenShift Container Platform 4.10 and later.



NOTE

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

6.3.4.3.2. Tested instance types for AWS

The following Amazon Web Services (AWS) instance types have been tested with OpenShift Container Platform.



NOTE

Use the machine types included in the following charts for your AWS instances. If you use an instance type that is not listed in the chart, ensure that the instance size you use matches the minimum resource requirements that are listed in the section named "Minimum resource requirements for cluster installation".

Example 6.22. Machine types based on 64-bit x86 architecture

- **c4.***
- **c5.***
- **c5a.***
- **i3.***
- **m4.***

- **m5.***
- **m5a.***
- **m6a.***
- **m6i.***
- **r4.***
- **r5.***
- **r5a.***
- **r6i.***
- **t3.***
- **t3a.***

6.3.4.3.3. Tested instance types for AWS on 64-bit ARM infrastructures

The following Amazon Web Services (AWS) 64-bit ARM instance types have been tested with OpenShift Container Platform.



NOTE

Use the machine types included in the following charts for your AWS ARM instances. If you use an instance type that is not listed in the chart, ensure that the instance size you use matches the minimum resource requirements that are listed in "Minimum resource requirements for cluster installation".

Example 6.23. Machine types based on 64-bit ARM architecture

- **c6g.***
- **m6g.***

6.3.4.3.4. Sample customized install-config.yaml file for AWS

You can customize the installation configuration file (**install-config.yaml**) to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.



IMPORTANT

This sample YAML file is provided for reference only. You must obtain your **install-config.yaml** file by using the installation program and modify it.

```
apiVersion: v1
baseDomain: example.com 1
credentialsMode: Mint 2
```

```
controlPlane: 3 4
  hyperthreading: Enabled 5
  name: master
  platform:
    aws:
      zones:
        - us-west-2a
        - us-west-2b
      rootVolume:
        iops: 4000
        size: 500
        type: io1 6
      metadataService:
        authentication: Optional 7
        type: m6i.xlarge
      replicas: 3
compute: 8
- hyperthreading: Enabled 9
  name: worker
  platform:
    aws:
      rootVolume:
        iops: 2000
        size: 500
        type: io1 10
      metadataService:
        authentication: Optional 11
        type: c5.4xlarge
      zones:
        - us-west-2c
      replicas: 3
metadata:
  name: test-cluster 12
networking: 13
  clusterNetwork:
    - cidr: 10.128.0.0/14
      hostPrefix: 23
  machineNetwork:
    - cidr: 10.0.0.0/16
  networkType: OVNKubernetes 14
  serviceNetwork:
    - 172.30.0.0/16
platform:
  aws:
    region: us-west-2 15
    propagateUserTags: true 16
    userTags:
      adminContact: jdoe
      costCenter: 7536
    subnets: 17
    - subnet-1
    - subnet-2
    - subnet-3
  amiID: ami-0c5d3e03c0ab9b19a 18
```

```

serviceEndpoints: 19
  - name: ec2
    url: https://vpce-id.ec2.us-west-2.vpce.amazonaws.com
  hostedZone: Z3URY6TWQ91KVV 20
fips: false 21
sshKey: ssh-ed25519 AAAA... 22
pullSecret: '{"auths":{"<local_registry>":{"auth": "<credentials>","email": "you@example.com"}}}' 23
additionalTrustBundle: | 24
  -----BEGIN CERTIFICATE-----
  <MY_TRUSTED_CA_CERT>
  -----END CERTIFICATE-----
imageContentSources: 25
  - mirrors:
    - <local_registry>/<local_repository_name>/release
    source: quay.io/openshift-release-dev/ocp-release
  - mirrors:
    - <local_registry>/<local_repository_name>/release
    source: quay.io/openshift-release-dev/ocp-v4.0-art-dev

```


- 1 12 15 Required. The installation program prompts you for this value.
- 2 Optional: Add this parameter to force the Cloud Credential Operator (CCO) to use the specified mode. By default, the CCO uses the root credentials in the **kube-system** namespace to dynamically try to determine the capabilities of the credentials. For details about CCO modes, see the "About the Cloud Credential Operator" section in the *Authentication and authorization* guide.
- 3 8 13 16 If you do not provide these parameters and values, the installation program provides the default value.
- 4 The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.
- 5 9 Whether to enable or disable simultaneous multithreading, or **hyperthreading**. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores. You can disable it by setting the parameter value to **Disabled**. If you disable simultaneous multithreading in some cluster machines, you must disable it in all cluster machines.



IMPORTANT


If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance. Use larger instance types, such as **m4.2xlarge** or **m5.2xlarge**, for your machines if you disable simultaneous multithreading.

- 6 10 To configure faster storage for etcd, especially for larger clusters, set the storage type as **io1** and set **iops** to **2000**.
- 7 11 Whether to require the [Amazon EC2 Instance Metadata Service v2](#) (IMDSv2). To require IMDSv2, set the parameter value to **Required**. To allow the use of both IMDSv1 and IMDSv2, set the parameter value to **Optional**. If no value is specified, both IMDSv1 and IMDSv2 are allowed.

**NOTE**

The IMDS configuration for control plane machines that is set during cluster installation can only be changed by using the AWS CLI. The IMDS configuration for compute machines can be changed by using compute machine sets.

- 14 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 17 If you provide your own VPC, specify subnets for each availability zone that your cluster uses.
- 18 The ID of the AMI used to boot machines for the cluster. If set, the AMI must belong to the same region as the cluster.
- 19 The AWS service endpoints. Custom endpoints are required when installing to an unknown AWS region. The endpoint URL must use the **https** protocol and the host must trust the certificate.
- 20 The ID of your existing Route 53 private hosted zone. Providing an existing hosted zone requires that you supply your own VPC and the hosted zone is already associated with the VPC prior to installing your cluster. If undefined, the installation program creates a new hosted zone.
- 21 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.

**IMPORTANT**

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 22 You can optionally provide the **sshKey** value that you use to access the machines in your cluster.

**NOTE**

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- 23 For **<local_registry>**, specify the registry domain name, and optionally the port, that your mirror registry uses to serve content. For example **registry.example.com** or **registry.example.com:5000**. For **<credentials>**, specify the base64-encoded user name and password for your mirror registry.
- 24 Provide the contents of the certificate file that you used for your mirror registry.
- 25

Provide the **imageContentSources** section from the output of the command to mirror the repository.

6.3.4.3.5. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure


1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: ec2.<aws_region>.amazonaws.com,elasticsearch.<aws_region>.amazonaws.com,s3.<aws_region>.amazonaws.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all

destinations. If you have added the Amazon **EC2**, **Elastic Load Balancing**, and **S3** VPC endpoints to your VPC, you must add these endpoints to the **noProxy** field.

- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.



NOTE

The installation program does not support the proxy **readinessEndpoints** field.



NOTE

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.



NOTE

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

6.3.4.4. Alternatives to storing administrator-level secrets in the kube-system project

By default, administrator secrets are stored in the **kube-system** project. If you configured the **credentialsMode** parameter in the **install-config.yaml** file to **Manual**, you must use one of the following alternatives:

- To manage long-term cloud credentials manually, follow the procedure in [Manually creating long-term credentials](#).
- To implement short-term credentials that are managed outside the cluster for individual components, follow the procedures in [Configuring an AWS cluster to use short-term credentials](#).

6.3.4.4.1. Manually creating long-term credentials

The Cloud Credential Operator (CCO) can be put into manual mode prior to installation in environments where the cloud identity and access management (IAM) APIs are not reachable, or the administrator prefers not to store an administrator-level credential secret in the cluster **kube-system** namespace.

Procedure

1. If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```
apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...
```

2. If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

3. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/' {print $3})
```

4. Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included \1 \
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \2 \
  --to=<path_to_directory_for_credentials_requests> \3
```

- 1 The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2 Specify the location of the **install-config.yaml** file.
- 3 Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

This command creates a YAML file for each **CredentialsRequest** object.

Sample CredentialsRequest object

```
apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
spec:
```

```

providerSpec:
  apiVersion: cloudcredential.openshift.io/v1
  kind: AWSProviderSpec
  statementEntries:
  - effect: Allow
    action:
    - iam:GetUser
    - iam:GetUserPolicy
    - iam:ListAccessKeys
    resource: "*"
  ...

```

5. Create YAML files for secrets in the **openshift-install** manifests directory that you generated previously. The secrets must be stored using the namespace and secret name defined in the **spec.secretRef** for each **CredentialsRequest** object.

Sample CredentialsRequest object with secrets

```

apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: AWSProviderSpec
    statementEntries:
    - effect: Allow
      action:
      - s3:CreateBucket
      - s3>DeleteBucket
      resource: "*"
    ...
  secretRef:
    name: <component_secret>
    namespace: <component_namespace>
  ...

```

Sample Secret object

```

apiVersion: v1
kind: Secret
metadata:
  name: <component_secret>
  namespace: <component_namespace>
data:
  aws_access_key_id: <base64_encoded_aws_access_key_id>
  aws_secret_access_key: <base64_encoded_aws_secret_access_key>

```




IMPORTANT

Before upgrading a cluster that uses manually maintained credentials, you must ensure that the CCO is in an upgradeable state.

6.3.4.4.2. Configuring an AWS cluster to use short-term credentials

To install a cluster that is configured to use the AWS Security Token Service (STS), you must configure the CCO utility and create the required AWS resources for your cluster.

6.3.4.4.2.1. Configuring the Cloud Credential Operator utility

To create and manage cloud credentials from outside of the cluster when the Cloud Credential Operator (CCO) is operating in manual mode, extract and prepare the CCO utility (**ccoctl**) binary.



NOTE

The **ccoctl** utility is a Linux binary that must run in a Linux environment.

Prerequisites

- You have access to an OpenShift Container Platform account with cluster administrator access.
- You have installed the OpenShift CLI (**oc**).
- You have created an AWS account for the **ccoctl** utility to use with the following permissions:

Example 6.24. Required AWS permissions

Required iam permissions

- **iam:CreateOpenIDConnectProvider**
- **iam:CreateRole**
- **iam>DeleteOpenIDConnectProvider**
- **iam>DeleteRole**
- **iam>DeleteRolePolicy**
- **iam:GetOpenIDConnectProvider**
- **iam:GetRole**
- **iam:GetUser**
- **iam:ListOpenIDConnectProviders**
- **iam:ListRolePolicies**
- **iam:ListRoles**
- **iam:PutRolePolicy**
- **iam:TagOpenIDConnectProvider**

- **iam:TagRole**

Required s3 permissions

- **s3:CreateBucket**
- **s3>DeleteBucket**
- **s3>DeleteObject**
- **s3:GetBucketAcl**
- **s3:GetBucketTagging**
- **s3:GetObject**
- **s3:GetObjectAcl**
- **s3:GetObjectTagging**
- **s3:ListBucket**
- **s3:PutBucketAcl**
- **s3:PutBucketPolicy**
- **s3:PutBucketPublicAccessBlock**
- **s3:PutBucketTagging**
- **s3:PutObject**
- **s3:PutObjectAcl**
- **s3:PutObjectTagging**

Required cloudfront permissions

- **cloudfront:ListCloudFrontOriginAccessIdentities**
- **cloudfront:ListDistributions**
- **cloudfront:ListTagsForResource**

If you plan to store the OIDC configuration in a private S3 bucket that is accessed by the IAM identity provider through a public CloudFront distribution URL, the AWS account that runs the **ccoctl** utility requires the following additional permissions:

Example 6.25. Additional permissions for a private S3 bucket with CloudFront

- **cloudfront:CreateCloudFrontOriginAccessIdentity**
- **cloudfront:CreateDistribution**
- **cloudfront>DeleteCloudFrontOriginAccessIdentity**
- **cloudfront>DeleteDistribution**

- **cloudfront:GetCloudFrontOriginAccessIdentity**
- **cloudfront:GetCloudFrontOriginAccessIdentityConfig**
- **cloudfront:GetDistribution**
- **cloudfront:TagResource**
- **cloudfront:UpdateDistribution**

**NOTE**

These additional permissions support the use of the **--create-private-s3-bucket** option when processing credentials requests with the **ccoctl aws create-all** command.


Procedure

1. Set a variable for the OpenShift Container Platform release image by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

2. Obtain the CCO container image from the OpenShift Container Platform release image by running the following command:

```
$ CCO_IMAGE=$(oc adm release info --image-for='cloud-credential-operator' $RELEASE_IMAGE -a ~/.pull-secret)
```

**NOTE**

Ensure that the architecture of the **\$RELEASE_IMAGE** matches the architecture of the environment in which you will use the **ccoctl** tool.

3. Extract the **ccoctl** binary from the CCO container image within the OpenShift Container Platform release image by running the following command:

```
$ oc image extract $CCO_IMAGE \
  --file="/usr/bin/ccoctl.<rhel_version>" 1 \
  -a ~/.pull-secret
```

- 1** For **<rhel_version>**, specify the value that corresponds to the version of Red Hat Enterprise Linux (RHEL) that the host uses. If no value is specified, **ccoctl.rhel8** is used by default. The following values are valid:

- **rhel8**: Specify this value for hosts that use RHEL 8.
- **rhel9**: Specify this value for hosts that use RHEL 9.

4. Change the permissions to make **ccoctl** executable by running the following command:

```
$ chmod 775 ccoctl.<rhel_version>
```

Verification

- To verify that **ccoctl** is ready to use, display the help file by running the following command:

```
$ ccoctl --help
```

Output of **ccoctl --help**

```
OpenShift credentials provisioning tool
```

```
Usage:
```

```
ccoctl [command]
```

```
Available Commands:
```

```
aws      Manage credentials objects for AWS cloud
azure    Manage credentials objects for Azure
gcp      Manage credentials objects for Google cloud
help     Help about any command
ibmcloud Manage credentials objects for IBM Cloud
nutanix  Manage credentials objects for Nutanix
```

```
Flags:
```

```
-h, --help  help for ccoctl
```

```
Use "ccoctl [command] --help" for more information about a command.
```

6.3.4.4.2.2. Creating AWS resources with the Cloud Credential Operator utility

You have the following options when creating AWS resources:

- You can use the **ccoctl aws create-all** command to create the AWS resources automatically. This is the quickest way to create the resources. See [Creating AWS resources with a single command](#).
- If you need to review the JSON files that the **ccoctl** tool creates before modifying AWS resources, or if the process the **ccoctl** tool uses to create AWS resources automatically does not meet the requirements of your organization, you can create the AWS resources individually. See [Creating AWS resources individually](#).

6.3.4.4.2.2.1. Creating AWS resources with a single command

If the process the **ccoctl** tool uses to create AWS resources automatically meets the requirements of your organization, you can use the **ccoctl aws create-all** command to automate the creation of AWS resources.

Otherwise, you can create the AWS resources individually. For more information, see "Creating AWS resources individually".



NOTE

By default, **ccoctl** creates objects in the directory in which the commands are run. To create the objects in a different directory, use the **--output-dir** flag. This procedure uses **<path_to_ccoctl_output_dir>** to refer to this directory.

Prerequisites

You must have:

- Extracted and prepared the **ccoctl** binary.

Procedure


1. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

2. Extract the list of **CredentialsRequest** objects from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included \1 \
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \2 \
  --to=<path_to_directory_for_credentials_requests> \3
```

- 1 The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2 Specify the location of the **install-config.yaml** file.
- 3 Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.



NOTE

This command might take a few moments to run.

3. Use the **ccoctl** tool to process all **CredentialsRequest** objects by running the following command:

```
$ ccoctl aws create-all \
  --name=<name> \1 \
  --region=<aws_region> \2 \
  --credentials-requests-dir=<path_to_credentials_requests_directory> \3 \
  --output-dir=<path_to_ccoctl_output_dir> \4 \
  --create-private-s3-bucket \5
```

- 1 Specify the name used to tag any cloud resources that are created for tracking.

- 2 Specify the AWS region in which cloud resources will be created.
- 3 Specify the directory containing the files for the component **CredentialsRequest** objects.
- 4 Optional: Specify the directory in which you want the **ccoctl** utility to create objects. By default, the utility creates objects in the directory in which the commands are run.
- 5 Optional: By default, the **ccoctl** utility stores the OpenID Connect (OIDC) configuration files in a public S3 bucket and uses the S3 URL as the public OIDC endpoint. To store the



NOTE

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

Verification

- To verify that the OpenShift Container Platform secrets are created, list the files in the `<path_to_ccoctl_output_dir>/manifests` directory:

```
$ ls <path_to_ccoctl_output_dir>/manifests
```

Example output

```
cluster-authentication-02-config.yaml
openshift-cloud-credential-operator-cloud-credential-operator-iam-ro-creds-credentials.yaml
openshift-cloud-network-config-controller-cloud-credentials-credentials.yaml
openshift-cluster-api-capamanager-bootstrap-credentials-credentials.yaml
openshift-cluster-csi-drivers-ebs-cloud-credentials-credentials.yaml
openshift-image-registry-installer-cloud-credentials-credentials.yaml
openshift-ingress-operator-cloud-credentials-credentials.yaml
openshift-machine-api-aws-cloud-credentials-credentials.yaml
```

You can verify that the IAM roles are created by querying AWS. For more information, refer to AWS documentation on listing IAM roles.

6.3.4.4.2.2.2. Creating AWS resources individually

You can use the **ccoctl** tool to create AWS resources individually. This option might be useful for an organization that shares the responsibility for creating these resources among different users or departments.

Otherwise, you can use the **ccoctl aws create-all** command to create the AWS resources automatically. For more information, see "Creating AWS resources with a single command".



NOTE

By default, **ccoctl** creates objects in the directory in which the commands are run. To create the objects in a different directory, use the **--output-dir** flag. This procedure uses **<path_to_ccoctl_output_dir>** to refer to this directory.

Some **ccoctl** commands make AWS API calls to create or modify AWS resources. You can use the **--dry-run** flag to avoid making API calls. Using this flag creates JSON files on the local file system instead. You can review and modify the JSON files and then apply them with the AWS CLI tool using the **--cli-input-json** parameters.

Prerequisites

- Extract and prepare the **ccoctl** binary.

Procedure

1. Generate the public and private RSA key files that are used to set up the OpenID Connect provider for the cluster by running the following command:

```
$ ccoctl aws create-key-pair
```

Example output

```
2021/04/13 11:01:02 Generating RSA keypair
2021/04/13 11:01:03 Writing private key to /<path_to_ccoctl_output_dir>/serviceaccount-
signer.private
2021/04/13 11:01:03 Writing public key to /<path_to_ccoctl_output_dir>/serviceaccount-
signer.public
2021/04/13 11:01:03 Copying signing key for use by installer
```

where **serviceaccount-signer.private** and **serviceaccount-signer.public** are the generated key files.

This command also creates a private key that the cluster requires during installation in **<path_to_ccoctl_output_dir>/tls/bound-service-account-signing-key.key**.

2. Create an OpenID Connect identity provider and S3 bucket on AWS by running the following command:

```
$ ccoctl aws create-identity-provider \
  --name=<name> \ 1
  --region=<aws_region> \ 2
  --public-key-file=<path_to_ccoctl_output_dir>/serviceaccount-signer.public 3
```

- 1** **<name>** is the name used to tag any cloud resources that are created for tracking.
- 2** **<aws-region>** is the AWS region in which cloud resources will be created.
- 3** **<path_to_ccoctl_output_dir>** is the path to the public key file that the **ccoctl aws create-key-pair** command generated.

Example output

-

```

2021/04/13 11:16:09 Bucket <name>-oidc created
2021/04/13 11:16:10 OpenID Connect discovery document in the S3 bucket <name>-oidc at
.well-known/openid-configuration updated
2021/04/13 11:16:10 Reading public key
2021/04/13 11:16:10 JSON web key set (JWKS) in the S3 bucket <name>-oidc at keys.json
updated
2021/04/13 11:16:18 Identity Provider created with ARN: arn:aws:iam::
<aws_account_id>:oidc-provider/<name>-oidc.s3.<aws_region>.amazonaws.com

```

where **openid-configuration** is a discovery document and **keys.json** is a JSON web key set file.

This command also creates a YAML configuration file in `/<path_to_ccoctl_output_dir>/manifests/cluster-authentication-02-config.yaml`. This file sets the issuer URL field for the service account tokens that the cluster generates, so that the AWS IAM identity provider trusts the tokens.

3. Create IAM roles for each component in the cluster:

- a. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

- b. Extract the list of **CredentialsRequest** objects from the OpenShift Container Platform release image:

```

$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included 1 \
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \
2 \
  --to=<path_to_directory_for_credentials_requests> 3

```

- 1** The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2** Specify the location of the **install-config.yaml** file.
- 3** Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

- c. Use the **ccoctl** tool to process all **CredentialsRequest** objects by running the following command:

```

$ ccoctl aws create-iam-roles \
  --name=<name> \
  --region=<aws_region> \
  --credentials-requests-dir=<path_to_credentials_requests_directory> \
  --identity-provider-arn=arn:aws:iam::<aws_account_id>:oidc-provider/<name>-oidc.s3.
<aws_region>.amazonaws.com

```




NOTE

For AWS environments that use alternative IAM API endpoints, such as GovCloud, you must also specify your region with the **--region** parameter.

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

For each **CredentialsRequest** object, **ccoctl** creates an IAM role with a trust policy that is tied to the specified OIDC identity provider, and a permissions policy as defined in each **CredentialsRequest** object from the OpenShift Container Platform release image.

Verification

- To verify that the OpenShift Container Platform secrets are created, list the files in the **<path_to_ccoctl_output_dir>/manifests** directory:

```
$ ls <path_to_ccoctl_output_dir>/manifests
```

Example output

```
cluster-authentication-02-config.yaml
openshift-cloud-credential-operator-cloud-credential-operator-iam-ro-creds-credentials.yaml
openshift-cloud-network-config-controller-cloud-credentials-credentials.yaml
openshift-cluster-api-capamanager-bootstrap-credentials-credentials.yaml
openshift-cluster-csi-drivers-ebs-cloud-credentials-credentials.yaml
openshift-image-registry-installer-cloud-credentials-credentials.yaml
openshift-ingress-operator-cloud-credentials-credentials.yaml
openshift-machine-api-aws-cloud-credentials-credentials.yaml
```

You can verify that the IAM roles are created by querying AWS. For more information, refer to AWS documentation on listing IAM roles.

6.3.4.4.2.3. Incorporating the Cloud Credential Operator utility manifests

To implement short-term security credentials managed outside the cluster for individual components, you must move the manifest files that the Cloud Credential Operator utility (**ccoctl**) created to the correct directories for the installation program.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have configured the Cloud Credential Operator utility (**ccoctl**).
- You have created the cloud provider resources that are required for your cluster with the **ccoctl** utility.

Procedure

1. If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```
apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...
```

- If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

- Copy the manifests that the **ccoctl** utility generated to the **manifests** directory that the installation program created by running the following command:

```
$ cp /<path_to_ccoctl_output_dir>/manifests/* ./manifests/
```

- Copy the private key that the **ccoctl** utility generated in the **tls** directory to the installation directory by running the following command:

```
$ cp -a /<path_to_ccoctl_output_dir>/tls .
```

6.3.4.5. Cluster Network Operator configuration

The configuration for the cluster network is specified as part of the Cluster Network Operator (CNO) configuration and stored in a custom resource (CR) object that is named **cluster**. The CR specifies the fields for the **Network** API in the **operator.openshift.io** API group.

The CNO configuration inherits the following fields during cluster installation from the **Network** API in the **Network.config.openshift.io** API group:

clusterNetwork

IP address pools from which pod IP addresses are allocated.

serviceNetwork

IP address pool for services.

defaultNetwork.type

Cluster network plugin. **OVNKubernetes** is the only supported plugin during installation.

You can specify the cluster network plugin configuration for your cluster by setting the fields for the **defaultNetwork** object in the CNO object named **cluster**.

6.3.4.5.1. Cluster Network Operator configuration object

The fields for the Cluster Network Operator (CNO) are described in the following table:

Table 6.3. Cluster Network Operator configuration object


Field	Type	Description
metadata.name	string	The name of the CNO object. This name is always cluster .
spec.clusterNetwork	array	A list specifying the blocks of IP addresses from which pod IP addresses are allocated and the subnet prefix length assigned to each individual node in the cluster. For example: <pre>spec: clusterNetwork: - cidr: 10.128.0.0/19 hostPrefix: 23 - cidr: 10.128.32.0/19 hostPrefix: 23</pre>
spec.serviceNetwork	array	A block of IP addresses for services. The OpenShift SDN and OVN-Kubernetes network plugins support only a single IP address block for the service network. For example: <pre>spec: serviceNetwork: - 172.30.0.0/14</pre> <p>You can customize this field only in the install-config.yaml file before you create the manifests. The value is read-only in the manifest file.</p>
spec.defaultNetwork	object	Configures the network plugin for the cluster network.
spec.kubeProxy Config	object	The fields for this object specify the kube-proxy configuration. If you are using the OVN-Kubernetes cluster network plugin, the kube-proxy configuration has no effect.

defaultNetwork object configuration

The values for the **defaultNetwork** object are defined in the following table:

Table 6.4. **defaultNetwork** object

Field	Type	Description
-------	------	-------------

Field	Type	Description
type	string	<p>OVNKubernetes. The Red Hat OpenShift Networking network plugin is selected during installation. This value cannot be changed after cluster installation.</p>  <p>NOTE</p> <p>OpenShift Container Platform uses the OVN-Kubernetes network plugin by default. OpenShift SDN is no longer available as an installation choice for new clusters.</p>
ovnKubernetesConfig	object	This object is only valid for the OVN-Kubernetes network plugin.

Configuration for the OVN-Kubernetes network plugin

The following table describes the configuration fields for the OVN-Kubernetes network plugin:

Table 6.5. ovnKubernetesConfig object

Field	Type	Description
mtu	integer	<p>The maximum transmission unit (MTU) for the Geneve (Generic Network Virtualization Encapsulation) overlay network. This is detected automatically based on the MTU of the primary network interface. You do not normally need to override the detected MTU.</p> <p>If the auto-detected value is not what you expect it to be, confirm that the MTU on the primary network interface on your nodes is correct. You cannot use this option to change the MTU value of the primary network interface on the nodes.</p> <p>If your cluster requires different MTU values for different nodes, you must set this value to 100 less than the lowest MTU value in your cluster. For example, if some nodes in your cluster have an MTU of 9001, and some have an MTU of 1500, you must set this value to 1400.</p>
genevePort	integer	The port to use for all Geneve packets. The default value is 6081 . This value cannot be changed after cluster installation.
ipsecConfig	object	Specify a configuration object for customizing the IPsec configuration.
ipv4	object	Specifies a configuration object for IPv4 settings.
ipv6	object	Specifies a configuration object for IPv6 settings.


Field	Type	Description
policyAuditConfig	object	Specify a configuration object for customizing network policy audit logging. If unset, the defaults audit log settings are used.
gatewayConfig	object	Optional: Specify a configuration object for customizing how egress traffic is sent to the node gateway.  <p>NOTE</p> <p>While migrating egress traffic, you can expect some disruption to workloads and service traffic until the Cluster Network Operator (CNO) successfully rolls out the changes.</p>

Table 6.6. `ovnKubernetesConfig.ipv4` object

Field	Type	Description
internalTransitSwitchSubnet	string	If your existing network infrastructure overlaps with the 100.88.0.0/16 IPv4 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. The subnet for the distributed transit switch that enables east-west traffic. This subnet cannot overlap with any other subnets used by OVN-Kubernetes or on the host itself. It must be large enough to accommodate one IP address per node in your cluster. The default value is 100.88.0.0/16 .
internalJoinSubnet	string	If your existing network infrastructure overlaps with the 100.64.0.0/16 IPv4 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. You must ensure that the IP address range does not overlap with any other subnet used by your OpenShift Container Platform installation. The IP address range must be larger than the maximum number of nodes that can be added to the cluster. For example, if the clusterNetwork.cidr value is 10.128.0.0/14 and the clusterNetwork.hostPrefix value is /23 , then the maximum number of nodes is 2⁽²³⁻¹⁴⁾=512 . The default value is 100.64.0.0/16 .

Table 6.7. `ovnKubernetesConfig.ipv6` object

Field	Type	Description
-------	------	-------------

Field	Type	Description
internalTransitSwitchSubnet	string	<p>If your existing network infrastructure overlaps with the fd97::/64 IPv6 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. The subnet for the distributed transit switch that enables east-west traffic. This subnet cannot overlap with any other subnets used by OVN-Kubernetes or on the host itself. It must be large enough to accommodate one IP address per node in your cluster.</p> <p>The default value is fd97::/64.</p>
internalJoinSubnet	string	<p>If your existing network infrastructure overlaps with the fd98::/64 IPv6 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. You must ensure that the IP address range does not overlap with any other subnet used by your OpenShift Container Platform installation. The IP address range must be larger than the maximum number of nodes that can be added to the cluster.</p> <p>The default value is fd98::/64.</p>

Table 6.8. **policyAuditConfig** object

Field	Type	Description
rateLimit	integer	The maximum number of messages to generate every second per node. The default value is 20 messages per second.
maxFileSize	integer	The maximum size for the audit log in bytes. The default value is 50000000 or 50 MB.
maxLogFiles	integer	The maximum number of log files that are retained.
destination	string	<p>One of the following additional audit log targets:</p> <p>libc The libc syslog() function of the journald process on the host.</p> <p>udp:<host>:<port> A syslog server. Replace <host>:<port> with the host and port of the syslog server.</p> <p>unix:<file> A Unix Domain Socket file specified by <file>.</p> <p>null Do not send the audit logs to any additional target.</p>
syslogFacility	string	The syslog facility, such as kern , as defined by RFC5424. The default value is local0 .

Table 6.9. gatewayConfig object

Field	Type	Description
routingViaHost	boolean	<p>Set this field to true to send egress traffic from pods to the host networking stack. For highly-specialized installations and applications that rely on manually configured routes in the kernel routing table, you might want to route egress traffic to the host networking stack. By default, egress traffic is processed in OVN to exit the cluster and is not affected by specialized routes in the kernel routing table. The default value is false.</p> <p>This field has an interaction with the Open vSwitch hardware offloading feature. If you set this field to true, you do not receive the performance benefits of the offloading because egress traffic is processed by the host networking stack.</p>
ipForwarding	object	<p>You can control IP forwarding for all traffic on OVN-Kubernetes managed interfaces by using the ipForwarding specification in the Network resource. Specify Restricted to only allow IP forwarding for Kubernetes related traffic. Specify Global to allow forwarding of all IP traffic. For new installations, the default is Restricted. For updates to OpenShift Container Platform 4.14 or later, the default is Global.</p>
ipv4	object	<p>Optional: Specify an object to configure the internal OVN-Kubernetes masquerade address for host to service traffic for IPv4 addresses.</p>
ipv6	object	<p>Optional: Specify an object to configure the internal OVN-Kubernetes masquerade address for host to service traffic for IPv6 addresses.</p>

Table 6.10. gatewayConfig.ipv4 object

Field	Type	Description
internalMasqueradeSubnet	string	<p>The masquerade IPv4 addresses that are used internally to enable host to service traffic. The host is configured with these IP addresses as well as the shared gateway bridge interface. The default value is 169.254.169.0/29.</p>

Table 6.11. gatewayConfig.ipv6 object

Field	Type	Description
internalMasqueradeSubnet	string	<p>The masquerade IPv6 addresses that are used internally to enable host to service traffic. The host is configured with these IP addresses as well as the shared gateway bridge interface. The default value is fd69::/125.</p>

Table 6.12. `ipsecConfig` object

Field	Type	Description
<code>mode</code>	<code>string</code>	<p>Specifies the behavior of the IPsec implementation. Must be one of the following values:</p> <ul style="list-style-type: none"> ● Disabled: IPsec is not enabled on cluster nodes. ● External: IPsec is enabled for network traffic with external hosts. ● Full: IPsec is enabled for pod traffic and network traffic with external hosts.

Example OVN-Kubernetes configuration with IPsec enabled

```
defaultNetwork:
  type: OVNKubernetes
  ovnKubernetesConfig:
    mtu: 1400
    genevePort: 6081
    ipsecConfig:
      mode: Full
```




IMPORTANT

Using OVNKubernetes can lead to a stack exhaustion problem on IBM Power®.

kubeProxyConfig object configuration (OpenShiftSDN container network interface only)

The values for the `kubeProxyConfig` object are defined in the following table:


Table 6.13. `kubeProxyConfig` object

Field	Type	Description
<code>iptablesSyncPeriod</code>	<code>string</code>	<p>The refresh period for <code>iptables</code> rules. The default value is 30s. Valid suffixes include s, m, and h and are described in the Go time package documentation.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p>NOTE</p> <p>Because of performance improvements introduced in OpenShift Container Platform 4.3 and greater, adjusting the <code>iptablesSyncPeriod</code> parameter is no longer necessary.</p> </div> </div>

Field	Type	Description
proxyArguments.iptables-min-sync-period	array	<p>The minimum duration before refreshing iptables rules. This field ensures that the refresh does not happen too frequently. Valid suffixes include s, m, and h and are described in the Go time package. The default value is:</p> <pre>kubeProxyConfig: proxyArguments: iptables-min-sync-period: - 0s</pre>

6.3.4.6. Specifying advanced network configuration

You can use advanced network configuration for your network plugin to integrate your cluster into your existing network environment. You can specify advanced network configuration only before you install the cluster.



IMPORTANT

Customizing your network configuration by modifying the OpenShift Container Platform manifest files created by the installation program is not supported. Applying a manifest file that you create, as in the following procedure, is supported.

Prerequisites

- You have created the **install-config.yaml** file and completed any modifications to it.

Procedure

- Change to the directory that contains the installation program and create the manifests:

```
$ ./openshift-install create manifests --dir <installation_directory> 1
```

- 1** **<installation_directory>** specifies the name of the directory that contains the **install-config.yaml** file for your cluster.

- Create a stub manifest file for the advanced network configuration that is named **cluster-network-03-config.yml** in the **<installation_directory>/manifests/** directory:

```
apiVersion: operator.openshift.io/v1
kind: Network
metadata:
  name: cluster
spec:
```

- Specify the advanced network configuration for your cluster in the **cluster-network-03-config.yml** file, such as in the following example:

Enable IPsec for the OVN-Kubernetes network provider

```

apiVersion: operator.openshift.io/v1
kind: Network
metadata:
  name: cluster
spec:
  defaultNetwork:
  ovnKubernetesConfig:
  ipsecConfig:
  mode: Full

```

- Optional: Back up the **manifests/cluster-network-03-config.yml** file. The installation program consumes the **manifests/** directory when you create the Ignition config files.



NOTE

For more information on using a Network Load Balancer (NLB) on AWS, see [Configuring Ingress cluster traffic on AWS using a Network Load Balancer](#).

6.3.4.7. Configuring an Ingress Controller Network Load Balancer on a new AWS cluster

You can create an Ingress Controller backed by an AWS Network Load Balancer (NLB) on a new cluster.

Prerequisites

- Create the **install-config.yaml** file and complete any modifications to it.

Procedure

Create an Ingress Controller backed by an AWS NLB on a new cluster.

- Change to the directory that contains the installation program and create the manifests:

```
$ ./openshift-install create manifests --dir <installation_directory> 1
```

- For **<installation_directory>**, specify the name of the directory that contains the **install-config.yaml** file for your cluster.

- Create a file that is named **cluster-ingress-default-ingresscontroller.yaml** in the **<installation_directory>/manifests/** directory:

```
$ touch <installation_directory>/manifests/cluster-ingress-default-ingresscontroller.yaml 1
```

- For **<installation_directory>**, specify the directory name that contains the **manifests/** directory for your cluster.

After creating the file, several network configuration files are in the **manifests/** directory, as shown:

```
$ ls <installation_directory>/manifests/cluster-ingress-default-ingresscontroller.yaml
```

Example output

```
cluster-ingress-default-ingresscontroller.yaml
```

- Open the **cluster-ingress-default-ingresscontroller.yaml** file in an editor and enter a custom resource (CR) that describes the Operator configuration you want:

```
apiVersion: operator.openshift.io/v1
kind: IngressController
metadata:
  creationTimestamp: null
  name: default
  namespace: openshift-ingress-operator
spec:
  endpointPublishingStrategy:
    loadBalancer:
      scope: External
      providerParameters:
        type: AWS
      aws:
        type: NLB
    type: LoadBalancerService
```

- Save the **cluster-ingress-default-ingresscontroller.yaml** file and quit the text editor.
- Optional: Back up the **manifests/cluster-ingress-default-ingresscontroller.yaml** file. The installation program deletes the **manifests/** directory when creating the cluster.

6.3.4.8. Configuring hybrid networking with OVN-Kubernetes

You can configure your cluster to use hybrid networking with the OVN-Kubernetes network plugin. This allows a hybrid cluster that supports different node networking configurations.

Prerequisites

- You defined **OVNKubernetes** for the **networking.networkType** parameter in the **install-config.yaml** file. See the installation documentation for configuring OpenShift Container Platform network customizations on your chosen cloud provider for more information.

Procedure

- Change to the directory that contains the installation program and create the manifests:

```
$ ./openshift-install create manifests --dir <installation_directory>
```

where:

<installation_directory>

Specifies the name of the directory that contains the **install-config.yaml** file for your cluster.

- Create a stub manifest file for the advanced network configuration that is named **cluster-network-03-config.yml** in the **<installation_directory>/manifests/** directory:

```
$ cat <<EOF > <installation_directory>/manifests/cluster-network-03-config.yml
apiVersion: operator.openshift.io/v1
```

```
kind: Network
metadata:
  name: cluster
spec:
EOF
```

where:

<installation_directory>

Specifies the directory name that contains the **manifests/** directory for your cluster.

- Open the **cluster-network-03-config.yml** file in an editor and configure OVN-Kubernetes with hybrid networking, such as in the following example:

Specify a hybrid networking configuration


```
apiVersion: operator.openshift.io/v1
kind: Network
metadata:
  name: cluster
spec:
  defaultNetwork:
    ovnKubernetesConfig:
      hybridOverlayConfig:
        hybridClusterNetwork: 1
        - cidr: 10.132.0.0/14
        hostPrefix: 23
```

- Specify the CIDR configuration used for nodes on the additional overlay network. The **hybridClusterNetwork** CIDR cannot overlap with the **clusterNetwork** CIDR.

- Save the **cluster-network-03-config.yml** file and quit the text editor.
- Optional: Back up the **manifests/cluster-network-03-config.yml** file. The installation program deletes the **manifests/** directory when creating the cluster.

6.3.4.9. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have the OpenShift Container Platform installation program and the pull secret for your cluster.

- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.

Procedure

- Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

- For **<installation_directory>**, specify the location of your customized **./install-config.yaml** file.
- To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

- Optional: Remove or disable the **AdministratorAccess** policy from the IAM account that you used to install the cluster.



NOTE

The elevated permissions provided by the **AdministratorAccess** policy are required only during installation.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to **<installation_directory>/openshift_install.log**.




IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

6.3.4.10. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

6.3.4.11. Logging in to the cluster by using the web console

The **kubeadmin** user exists by default after an OpenShift Container Platform installation. You can log in to your cluster as the **kubeadmin** user by using the OpenShift Container Platform web console.


Prerequisites

- You have access to the installation host.
- You completed a cluster installation and all cluster Operators are available.

Procedure

1. Obtain the password for the **kubeadmin** user from the **kubeadmin-password** file on the installation host:

```
$ cat <installation_directory>/auth/kubeadmin-password
```



NOTE

Alternatively, you can obtain the **kubeadmin** password from the **<installation_directory>/openshift_install.log** log file on the installation host.

2. List the OpenShift Container Platform web console route:

```
$ oc get routes -n openshift-console | grep 'console-openshift'
```



NOTE

Alternatively, you can obtain the OpenShift Container Platform route from the **<installation_directory>/openshift_install.log** log file on the installation host.

Example output

```
console    console-openshift-console.apps.<cluster_name>.<base_domain>    console
https reencrypt/Redirect None
```

3. Navigate to the route detailed in the output of the preceding command in a web browser and log in as the **kubeadmin** user.

Additional resources

- See [Accessing the web console](#) for more details about accessing and understanding the OpenShift Container Platform web console.

6.3.4.12. Next steps

- [Validating an installation.](#)
- [Customize your cluster.](#)
- If necessary, you can [opt out of remote health reporting](#) .
- If necessary, you can [remove cloud provider credentials](#).

6.3.5. Installing a cluster on AWS in a restricted network

In OpenShift Container Platform version 4.16, you can install a cluster on Amazon Web Services (AWS) in a restricted network by creating an internal mirror of the installation release content on an existing Amazon Virtual Private Cloud (VPC).

6.3.5.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [mirrored the images for a disconnected installation](#) to your registry and obtained the **imageContentSources** data for your version of OpenShift Container Platform.



IMPORTANT

Because the installation media is on the mirror host, you can use that computer to complete all installation steps.

- You have an existing VPC in AWS. When installing to a restricted network using installer-provisioned infrastructure, you cannot use the installer-provisioned VPC. You must use a user-provisioned VPC that satisfies one of the following requirements:
 - Contains the mirror registry
 - Has firewall rules or a peering connection to access the mirror registry hosted elsewhere
- You [configured an AWS account](#) to host the cluster.



IMPORTANT

If you have an AWS profile stored on your computer, it must not use a temporary session token that you generated while using a multi-factor authentication device. The cluster continues to use your current AWS credentials to create AWS resources for the entire life of the cluster, so you must use key-based, long-term credentials. To generate appropriate keys, see [Managing Access Keys for IAM Users](#) in the AWS documentation. You can supply the keys when you run the installation program.

- You downloaded the AWS CLI and installed it on your computer. See [Install the AWS CLI Using the Bundled Installer \(Linux, macOS, or UNIX\)](#) in the AWS documentation.
- If you use a firewall and plan to use the Telemetry service, you [configured the firewall to allow the sites](#) that your cluster requires access to.



NOTE

If you are configuring a proxy, be sure to also review this site list.

6.3.5.2. About installations in restricted networks

In OpenShift Container Platform 4.16, you can perform an installation that does not require an active connection to the internet to obtain software components. Restricted network installations can be

completed using installer-provisioned infrastructure or user-provisioned infrastructure, depending on the cloud platform to which you are installing the cluster.

If you choose to perform a restricted network installation on a cloud platform, you still require access to its cloud APIs. Some cloud functions, like Amazon Web Service's Route 53 DNS and IAM services, require internet access. Depending on your network, you might require less internet access for an installation on bare metal hardware, Nutanix, or on VMware vSphere.

To complete a restricted network installation, you must create a registry that mirrors the contents of the OpenShift image registry and contains the installation media. You can create this registry on a mirror host, which can access both the internet and your closed network, or by using other methods that meet your restrictions.

6.3.5.2.1. Additional limits

Clusters in restricted networks have the following additional limitations and restrictions:

- The **ClusterVersion** status includes an **Unable to retrieve available updates** error.
- By default, you cannot use the contents of the Developer Catalog because you cannot access the required image stream tags.

6.3.5.3. About using a custom VPC

In OpenShift Container Platform 4.16, you can deploy a cluster into existing subnets in an existing Amazon Virtual Private Cloud (VPC) in Amazon Web Services (AWS). By deploying OpenShift Container Platform into an existing AWS VPC, you might be able to avoid limit constraints in new accounts or more easily abide by the operational constraints that your company's guidelines set. If you cannot obtain the infrastructure creation permissions that are required to create the VPC yourself, use this installation option.

Because the installation program cannot know what other components are also in your existing subnets, it cannot choose subnet CIDRs and so forth on your behalf. You must configure networking for the subnets that you install your cluster to yourself.

6.3.5.3.1. Requirements for using your VPC

The installation program no longer creates the following components:

- Internet gateways
- NAT gateways
- Subnets
- Route tables
- VPCs
- VPC DHCP options
- VPC endpoints

**NOTE**

The installation program requires that you use the cloud-provided DNS server. Using a custom DNS server is not supported and causes the installation to fail.

If you use a custom VPC, you must correctly configure it and its subnets for the installation program and the cluster to use. See [Amazon VPC console wizard configurations](#) and [Work with VPCs and subnets](#) in the AWS documentation for more information on creating and managing an AWS VPC.

The installation program cannot:

- Subdivide network ranges for the cluster to use.
- Set route tables for the subnets.
- Set VPC options like DHCP.

You must complete these tasks before you install the cluster. See [VPC networking components](#) and [Route tables for your VPC](#) for more information on configuring networking in an AWS VPC.

Your VPC must meet the following characteristics:

- The VPC must not use the **kubernetes.io/cluster/.*: owned, Name**, and **openshift.io/cluster** tags.
The installation program modifies your subnets to add the **kubernetes.io/cluster/.*: shared** tag, so your subnets must have at least one free tag slot available for it. See [Tag Restrictions](#) in the AWS documentation to confirm that the installation program can add a tag to each subnet that you specify. You cannot use a **Name** tag, because it overlaps with the EC2 **Name** field and the installation fails.
- If you want to extend your OpenShift Container Platform cluster into an AWS Outpost and have an existing Outpost subnet, the existing subnet must use the **kubernetes.io/cluster/unmanaged: true** tag. If you do not apply this tag, the installation might fail due to the Cloud Controller Manager creating a service load balancer in the Outpost subnet, which is an unsupported configuration.
- You must enable the **enableDnsSupport** and **enableDnsHostnames** attributes in your VPC, so that the cluster can use the Route 53 zones that are attached to the VPC to resolve cluster's internal DNS records. See [DNS Support in Your VPC](#) in the AWS documentation.
If you prefer to use your own Route 53 hosted private zone, you must associate the existing hosted zone with your VPC prior to installing a cluster. You can define your hosted zone using the **platform.aws.hostedZone** and **platform.aws.hostedZoneRole** fields in the **install-config.yaml** file. You can use a private hosted zone from another account by sharing it with the account where you install the cluster. If you use a private hosted zone from another account, you must use the **Passthrough** or **Manual** credentials mode.

If you are working in a disconnected environment, you are unable to reach the public IP addresses for EC2, ELB, and S3 endpoints. Depending on the level to which you want to restrict internet traffic during the installation, the following configuration options are available:

Option 1: Create VPC endpoints

Create a VPC endpoint and attach it to the subnets that the clusters are using. Name the endpoints as follows:

- **ec2.<aws_region>.amazonaws.com**
- **elasticloadbalancing.<aws_region>.amazonaws.com**

- **s3.<aws_region>.amazonaws.com**

With this option, network traffic remains private between your VPC and the required AWS services.

Option 2: Create a proxy without VPC endpoints

As part of the installation process, you can configure an HTTP or HTTPS proxy. With this option, internet traffic goes through the proxy to reach the required AWS services.

Option 3: Create a proxy with VPC endpoints

As part of the installation process, you can configure an HTTP or HTTPS proxy with VPC endpoints. Create a VPC endpoint and attach it to the subnets that the clusters are using. Name the endpoints as follows:

- **ec2.<aws_region>.amazonaws.com**
- **elasticloadbalancing.<aws_region>.amazonaws.com**
- **s3.<aws_region>.amazonaws.com**

When configuring the proxy in the **install-config.yaml** file, add these endpoints to the **noProxy** field. With this option, the proxy prevents the cluster from accessing the internet directly. However, network traffic remains private between your VPC and the required AWS services.

Required VPC components

You must provide a suitable VPC and subnets that allow communication to your machines.

Component	AWS type	Description
VPC	<ul style="list-style-type: none"> • AWS::EC2::VPC • AWS::EC2::VPCEndpoint 	You must provide a public VPC for the cluster to use. The VPC uses an endpoint that references the route tables for each subnet to improve communication with the registry that is hosted in S3.
Public subnets	<ul style="list-style-type: none"> • AWS::EC2::Subnet • AWS::EC2::SubnetNetworkACLAssociation 	Your VPC must have public subnets for between 1 and 3 availability zones and associate them with appropriate Ingress rules.

Component	AWS type	Description												
Internet gateway	<ul style="list-style-type: none"> ● AWS::EC2::InternetGateway ● AWS::EC2::VPCElasticGatewayAttachment ● AWS::EC2::RouteTable ● AWS::EC2::Route ● AWS::EC2::SubnetRouteTableAssociation ● AWS::EC2::NatGateway ● AWS::EC2::EIP 	<p>You must have a public internet gateway, with public routes, attached to the VPC. In the provided templates, each public subnet has a NAT gateway with an EIP address. These NAT gateways allow cluster resources, like private subnet instances, to reach the internet and are not required for some restricted network or proxy scenarios.</p>												
Network access control	<ul style="list-style-type: none"> ● AWS::EC2::NetworkAcl ● AWS::EC2::NetworkAclEntry 	<p>You must allow the VPC to access the following ports:</p> <table border="1"> <thead> <tr> <th>Port</th> <th>Reason</th> </tr> </thead> <tbody> <tr> <td>80</td> <td>Inbound HTTP traffic</td> </tr> <tr> <td>443</td> <td>Inbound HTTPS traffic</td> </tr> <tr> <td>22</td> <td>Inbound SSH traffic</td> </tr> <tr> <td>1024 - 65535</td> <td>Inbound ephemeral traffic</td> </tr> <tr> <td>0 - 65535</td> <td>Outbound ephemeral traffic</td> </tr> </tbody> </table>	Port	Reason	80	Inbound HTTP traffic	443	Inbound HTTPS traffic	22	Inbound SSH traffic	1024 - 65535	Inbound ephemeral traffic	0 - 65535	Outbound ephemeral traffic
Port	Reason													
80	Inbound HTTP traffic													
443	Inbound HTTPS traffic													
22	Inbound SSH traffic													
1024 - 65535	Inbound ephemeral traffic													
0 - 65535	Outbound ephemeral traffic													
Private subnets	<ul style="list-style-type: none"> ● AWS::EC2::Subnet ● AWS::EC2::RouteTable ● AWS::EC2::SubnetRouteTableAssociation 	<p>Your VPC can have private subnets. The provided CloudFormation templates can create private subnets for between 1 and 3 availability zones. If you use private subnets, you must provide appropriate routes and tables for them.</p>												

6.3.5.3.2. VPC validation

To ensure that the subnets that you provide are suitable, the installation program confirms the following data:

- All the subnets that you specify exist.
- You provide private subnets.
- The subnet CIDRs belong to the machine CIDR that you specified.
- You provide subnets for each availability zone. Each availability zone contains no more than one public and one private subnet. If you use a private cluster, provide only a private subnet for each availability zone. Otherwise, provide exactly one public and private subnet for each availability zone.
- You provide a public subnet for each private subnet availability zone. Machines are not provisioned in availability zones that you do not provide private subnets for.

If you destroy a cluster that uses an existing VPC, the VPC is not deleted. When you remove the OpenShift Container Platform cluster from a VPC, the **kubernetes.io/cluster/.*: shared** tag is removed from the subnets that it used.

6.3.5.3.3. Division of permissions

Starting with OpenShift Container Platform 4.3, you do not need all of the permissions that are required for an installation program-provisioned infrastructure cluster to deploy a cluster. This change mimics the division of permissions that you might have at your company: some individuals can create different resource in your clouds than others. For example, you might be able to create application-specific items, like instances, buckets, and load balancers, but not networking-related components such as VPCs, subnets, or ingress rules.

The AWS credentials that you use when you create your cluster do not need the networking permissions that are required to make VPCs and core networking components within the VPC, such as subnets, routing tables, internet gateways, NAT, and VPN. You still need permission to make the application resources that the machines within the cluster require, such as ELBs, security groups, S3 buckets, and nodes.

6.3.5.3.4. Isolation between clusters

If you deploy OpenShift Container Platform to an existing network, the isolation of cluster services is reduced in the following ways:

- You can install multiple OpenShift Container Platform clusters in the same VPC.
- ICMP ingress is allowed from the entire network.
- TCP 22 ingress (SSH) is allowed to the entire network.
- Control plane TCP 6443 ingress (Kubernetes API) is allowed to the entire network.
- Control plane TCP 22623 ingress (MCS) is allowed to the entire network.

6.3.5.4. Creating the installation configuration file

You can customize the OpenShift Container Platform cluster you install on Amazon Web Services (AWS).

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster. For a restricted network installation, these files are on your mirror host.
- You have the **imageContentSources** values that were generated during mirror registry creation.
- You have obtained the contents of the certificate for your mirror registry.

Procedure

1. Create the **install-config.yaml** file.
 - a. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install create install-config --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the directory name to store the files that the installation program creates.

When specifying the directory:

- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
- Use an empty directory. Some installation assets, such as bootstrap X.509 certificates, have short expiration intervals, therefore you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

- b. At the prompts, provide the configuration details for your cloud:
 - i. Optional: Select an SSH key to use to access your cluster machines.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- ii. Select **AWS** as the platform to target.
- iii. If you do not have an Amazon Web Services (AWS) profile stored on your computer, enter the AWS access key ID and secret access key for the user that you configured to run the installation program.
- iv. Select the AWS region to deploy the cluster to.
- v. Select the base domain for the Route 53 service that you configured for your cluster.
- vi. Enter a descriptive name for your cluster.

2. Edit the **install-config.yaml** file to give the additional information that is required for an installation in a restricted network.

- a. Update the **pullSecret** value to contain the authentication information for your registry:

```
pullSecret: '{"auths":{"<mirror_host_name>:5000":{"auth": "<credentials>","email":
"you@example.com"}}}'
```

For **<mirror_host_name>**, specify the registry domain name that you specified in the certificate for your mirror registry, and for **<credentials>**, specify the base64-encoded user name and password for your mirror registry.

- b. Add the **additionalTrustBundle** parameter and value.

```
additionalTrustBundle: |
  -----BEGIN CERTIFICATE-----

  ////////////////////////////////////////////////////////////////////
  -----END CERTIFICATE-----
```

The value must be the contents of the certificate file that you used for your mirror registry. The certificate file can be an existing, trusted certificate authority, or the self-signed certificate that you generated for the mirror registry.

- c. Define the subnets for the VPC to install the cluster in:

```
subnets:
- subnet-1
- subnet-2
- subnet-3
```

- d. Add the image content resources, which resemble the following YAML excerpt:

```
imageContentSources:
- mirrors:
  - <mirror_host_name>:5000/<repo_name>/release
  source: quay.io/openshift-release-dev/ocp-release
- mirrors:
  - <mirror_host_name>:5000/<repo_name>/release
  source: registry.redhat.io/ocp/release
```


For these values, use the **imageContentSources** that you recorded during mirror registry creation.

- e. Optional: Set the publishing strategy to **Internal**:

```
publish: Internal
```

By setting this option, you create an internal Ingress Controller and a private load balancer.

3. Make any other modifications to the **install-config.yaml** file that you require. For more information about the parameters, see "Installation configuration parameters".
4. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

Additional resources

- [Installation configuration parameters for AWS](#)

6.3.5.4.1. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 6.14. Minimum resource requirements

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Bootstrap	RHCOS	4	16 GB	100 GB	300
Control plane	RHCOS	4	16 GB	100 GB	300
Compute	RHCOS, RHEL 8.6 and later [3]	2	8 GB	100 GB	300

1. One vCPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio: $(\text{threads per core} \times \text{cores}) \times \text{sockets} = \text{vCPUs}$.
2. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes which require a 10 ms p99 fsync duration. Note that on many cloud platforms, storage size and IOPS scale together, so you might need to over-allocate storage volume to obtain sufficient performance.
3. As with all user-provisioned installations, if you choose to use RHEL compute machines in your cluster, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Use of RHEL 7 compute machines is deprecated and has been removed in OpenShift Container Platform 4.10 and later.



NOTE

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).


If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

6.3.5.4.2. Sample customized install-config.yaml file for AWS

You can customize the installation configuration file (**install-config.yaml**) to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.



IMPORTANT

This sample YAML file is provided for reference only. You must obtain your **install-config.yaml** file by using the installation program and modify it.

```

apiVersion: v1
baseDomain: example.com 1
credentialsMode: Mint 2
controlPlane: 3 4
  hyperthreading: Enabled 5
name: master
platform:
  aws:
    zones:
    - us-west-2a
    - us-west-2b
  rootVolume:
    iops: 4000
    size: 500
    type: io1 6
  metadataService:
    authentication: Optional 7
    type: m6i.xlarge
  replicas: 3
compute: 8
  - hyperthreading: Enabled 9

```

```

name: worker
platform:
  aws:
    rootVolume:
      iops: 2000
      size: 500
      type: io1 10
    metadataService:
      authentication: Optional 11
      type: c5.4xlarge
      zones:
        - us-west-2c
    replicas: 3
metadata:
  name: test-cluster 12
networking:
  clusterNetwork:
    - cidr: 10.128.0.0/14
      hostPrefix: 23
  machineNetwork:
    - cidr: 10.0.0.0/16
  networkType: OVNKubernetes 13
  serviceNetwork:
    - 172.30.0.0/16
platform:
  aws:
    region: us-west-2 14
    propagateUserTags: true 15
    userTags:
      adminContact: jdoe
      costCenter: 7536
    subnets: 16
    - subnet-1
    - subnet-2
    - subnet-3
    amiID: ami-0c5d3e03c0ab9b19a 17
    serviceEndpoints: 18
    - name: ec2
      url: https://vpce-id.ec2.us-west-2.vpce.amazonaws.com
    hostedZone: Z3URY6TWQ91KVV 19
  fips: false 20
  sshKey: ssh-ed25519 AAAA... 21
  pullSecret: '{"auths":{"<local_registry>":{"auth":"<credentials>","email":"you@example.com"}}}' 22
  additionalTrustBundle: | 23
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  imageContentSources: 24
  - mirrors:
    - <local_registry>/<local_repository_name>/release
      source: quay.io/openshift-release-dev/ocp-release
  - mirrors:
    - <local_registry>/<local_repository_name>/release
      source: quay.io/openshift-release-dev/ocp-v4.0-art-dev

```

- 1 12 14 Required. The installation program prompts you for this value.
- 2 Optional: Add this parameter to force the Cloud Credential Operator (CCO) to use the specified mode. By default, the CCO uses the root credentials in the **kube-system** namespace to dynamically try to determine the capabilities of the credentials. For details about CCO modes, see the "About the Cloud Credential Operator" section in the *Authentication and authorization* guide.
- 3 8 15 If you do not provide these parameters and values, the installation program provides the default value.
- 4 The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.
- 5 9 Whether to enable or disable simultaneous multithreading, or **hyperthreading**. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores. You can disable it by setting the parameter value to **Disabled**. If you disable simultaneous multithreading in some cluster machines, you must disable it in all cluster machines.



IMPORTANT

If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance. Use larger instance types, such as **m4.2xlarge** or **m5.2xlarge**, for your machines if you disable simultaneous multithreading.

- 6 10 To configure faster storage for etcd, especially for larger clusters, set the storage type as **io1** and set **iopts** to **2000**.
- 7 11 Whether to require the [Amazon EC2 Instance Metadata Service v2](#) (IMDSv2). To require IMDSv2, set the parameter value to **Required**. To allow the use of both IMDSv1 and IMDSv2, set the parameter value to **Optional**. If no value is specified, both IMDSv1 and IMDSv2 are allowed.



NOTE

The IMDS configuration for control plane machines that is set during cluster installation can only be changed by using the AWS CLI. The IMDS configuration for compute machines can be changed by using compute machine sets.

- 13 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 16 If you provide your own VPC, specify subnets for each availability zone that your cluster uses.
- 17 The ID of the AMI used to boot machines for the cluster. If set, the AMI must belong to the same region as the cluster.
- 18 The AWS service endpoints. Custom endpoints are required when installing to an unknown AWS region. The endpoint URL must use the **https** protocol and the host must trust the certificate.
- 19 The ID of your existing Route 53 private hosted zone. Providing an existing hosted zone requires that you supply your own VPC and the hosted zone is already associated with the VPC prior to installing your cluster. If undefined, the installation program creates a new hosted zone.

- 20 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container



IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 21 You can optionally provide the **sshKey** value that you use to access the machines in your cluster.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.


- 22 For **<local_registry>**, specify the registry domain name, and optionally the port, that your mirror registry uses to serve content. For example **registry.example.com** or **registry.example.com:5000**. For **<credentials>**, specify the base64-encoded user name and password for your mirror registry.
- 23 Provide the contents of the certificate file that you used for your mirror registry.
- 24 Provide the **imageContentSources** section from the output of the command to mirror the repository.

6.3.5.4.3. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.


For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure


1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: ec2.<aws_region>.amazonaws.com,elasticloadbalancing.
<aws_region>.amazonaws.com,s3.<aws_region>.amazonaws.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations. If you have added the Amazon **EC2**, **Elastic Load Balancing**, and **S3** VPC endpoints to your VPC, you must add these endpoints to the **noProxy** field.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.

**NOTE**

The installation program does not support the proxy **readinessEndpoints** field.

**NOTE**

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.

**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

6.3.5.5. Alternatives to storing administrator-level secrets in the kube-system project

By default, administrator secrets are stored in the **kube-system** project. If you configured the **credentialsMode** parameter in the **install-config.yaml** file to **Manual**, you must use one of the following alternatives:

- To manage long-term cloud credentials manually, follow the procedure in [Manually creating long-term credentials](#).
- To implement short-term credentials that are managed outside the cluster for individual components, follow the procedures in [Configuring an AWS cluster to use short-term credentials](#).

6.3.5.5.1. Manually creating long-term credentials

The Cloud Credential Operator (CCO) can be put into manual mode prior to installation in environments where the cloud identity and access management (IAM) APIs are not reachable, or the administrator prefers not to store an administrator-level credential secret in the cluster **kube-system** namespace.

Procedure

1. If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```
apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...
```

- If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

- Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

- Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included \1
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \2
  --to=<path_to_directory_for_credentials_requests> \3
```

- The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- Specify the location of the **install-config.yaml** file.
- Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

This command creates a YAML file for each **CredentialsRequest** object.

Sample **CredentialsRequest** object

```
apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: AWSProviderSpec
    statementEntries:
      - effect: Allow
        action:
          - iam:GetUser
          - iam:GetUserPolicy
          - iam:ListAccessKeys
        resource: "*"
  ...
```

5. Create YAML files for secrets in the **openshift-install** manifests directory that you generated previously. The secrets must be stored using the namespace and secret name defined in the **spec.secretRef** for each **CredentialsRequest** object.

Sample CredentialsRequest object with secrets

```

apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: AWSProviderSpec
    statementEntries:
    - effect: Allow
      action:
      - s3:CreateBucket
      - s3>DeleteBucket
      resource: "*"
    ...
  secretRef:
    name: <component_secret>
    namespace: <component_namespace>
  ...

```

Sample Secret object

```

apiVersion: v1
kind: Secret
metadata:
  name: <component_secret>
  namespace: <component_namespace>
data:
  aws_access_key_id: <base64_encoded_aws_access_key_id>
  aws_secret_access_key: <base64_encoded_aws_secret_access_key>

```



IMPORTANT

Before upgrading a cluster that uses manually maintained credentials, you must ensure that the CCO is in an upgradeable state.

6.3.5.5.2. Configuring an AWS cluster to use short-term credentials

To install a cluster that is configured to use the AWS Security Token Service (STS), you must configure the CCO utility and create the required AWS resources for your cluster.

6.3.5.5.2.1. Configuring the Cloud Credential Operator utility

To create and manage cloud credentials from outside of the cluster when the Cloud Credential Operator (CCO) is operating in manual mode, extract and prepare the CCO utility (**ccoctl**) binary.

**NOTE**

The **ccoctl** utility is a Linux binary that must run in a Linux environment.

Prerequisites

- You have access to an OpenShift Container Platform account with cluster administrator access.
- You have installed the OpenShift CLI (**oc**).
- You have created an AWS account for the **ccoctl** utility to use with the following permissions:

Example 6.26. Required AWS permissions**Required iam permissions**

- **iam:CreateOpenIDConnectProvider**
- **iam:CreateRole**
- **iam>DeleteOpenIDConnectProvider**
- **iam>DeleteRole**
- **iam>DeleteRolePolicy**
- **iam:GetOpenIDConnectProvider**
- **iam:GetRole**
- **iam:GetUser**
- **iam:ListOpenIDConnectProviders**
- **iam:ListRolePolicies**
- **iam:ListRoles**
- **iam:PutRolePolicy**
- **iam:TagOpenIDConnectProvider**
- **iam:TagRole**

Required s3 permissions

- **s3:CreateBucket**
- **s3>DeleteBucket**
- **s3>DeleteObject**
- **s3:GetBucketAcl**
- **s3:GetBucketTagging**
- **s3:GetObject**
- **s3:GetObjectAcl**

- **s3:GetObjectTagging**
- **s3:ListBucket**
- **s3:PutBucketAcl**
- **s3:PutBucketPolicy**
- **s3:PutBucketPublicAccessBlock**
- **s3:PutBucketTagging**
- **s3:PutObject**
- **s3:PutObjectAcl**
- **s3:PutObjectTagging**

Required **cloudfront** permissions

- **cloudfront:ListCloudFrontOriginAccessIdentities**
- **cloudfront:ListDistributions**
- **cloudfront:ListTagsForResource**

If you plan to store the OIDC configuration in a private S3 bucket that is accessed by the IAM identity provider through a public CloudFront distribution URL, the AWS account that runs the **ccoctl** utility requires the following additional permissions:

Example 6.27. Additional permissions for a private S3 bucket with CloudFront

- **cloudfront:CreateCloudFrontOriginAccessIdentity**
- **cloudfront:CreateDistribution**
- **cloudfront>DeleteCloudFrontOriginAccessIdentity**
- **cloudfront>DeleteDistribution**
- **cloudfront:GetCloudFrontOriginAccessIdentity**
- **cloudfront:GetCloudFrontOriginAccessIdentityConfig**
- **cloudfront:GetDistribution**
- **cloudfront:TagResource**
- **cloudfront:UpdateDistribution**



NOTE

These additional permissions support the use of the **--create-private-s3-bucket** option when processing credentials requests with the **ccoctl aws create-all** command.

■

Procedure

1. Set a variable for the OpenShift Container Platform release image by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

2. Obtain the CCO container image from the OpenShift Container Platform release image by running the following command:

```
$ CCO_IMAGE=$(oc adm release info --image-for='cloud-credential-operator'
$RELEASE_IMAGE -a ~/.pull-secret)
```



NOTE

Ensure that the architecture of the **\$RELEASE_IMAGE** matches the architecture of the environment in which you will use the **ccoctl** tool.

3. Extract the **ccoctl** binary from the CCO container image within the OpenShift Container Platform release image by running the following command:

```
$ oc image extract $CCO_IMAGE \
--file="/usr/bin/ccoctl.<rhel_version>" \
-a ~/.pull-secret
```

- 1 For **<rhel_version>**, specify the value that corresponds to the version of Red Hat Enterprise Linux (RHEL) that the host uses. If no value is specified, **ccoctl.rhel8** is used by default. The following values are valid:

- **rhel8**: Specify this value for hosts that use RHEL 8.
- **rhel9**: Specify this value for hosts that use RHEL 9.

4. Change the permissions to make **ccoctl** executable by running the following command:

```
$ chmod 775 ccoctl.<rhel_version>
```

Verification

- To verify that **ccoctl** is ready to use, display the help file by running the following command:

```
$ ccoctl --help
```

Output of **ccoctl --help**

```
OpenShift credentials provisioning tool
```

```
Usage:
ccoctl [command]
```

Available Commands:

aws Manage credentials objects for AWS cloud
 azure Manage credentials objects for Azure
 gcp Manage credentials objects for Google cloud
 help Help about any command
 ibmcloud Manage credentials objects for IBM Cloud
 nutanix Manage credentials objects for Nutanix

Flags:

-h, --help help for ccoctl

Use "ccoctl [command] --help" for more information about a command.

6.3.5.5.2.2. Creating AWS resources with the Cloud Credential Operator utility

You have the following options when creating AWS resources:

- You can use the **ccoctl aws create-all** command to create the AWS resources automatically. This is the quickest way to create the resources. See [Creating AWS resources with a single command](#).
- If you need to review the JSON files that the **ccoctl** tool creates before modifying AWS resources, or if the process the **ccoctl** tool uses to create AWS resources automatically does not meet the requirements of your organization, you can create the AWS resources individually. See [Creating AWS resources individually](#).

6.3.5.5.2.2.1. Creating AWS resources with a single command

If the process the **ccoctl** tool uses to create AWS resources automatically meets the requirements of your organization, you can use the **ccoctl aws create-all** command to automate the creation of AWS resources.

Otherwise, you can create the AWS resources individually. For more information, see "Creating AWS resources individually".

**NOTE**

By default, **ccoctl** creates objects in the directory in which the commands are run. To create the objects in a different directory, use the **--output-dir** flag. This procedure uses **<path_to_ccoctl_output_dir>** to refer to this directory.

Prerequisites

You must have:

- Extracted and prepared the **ccoctl** binary.

Procedure

1. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

2. Extract the list of **CredentialsRequest** objects from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included \ 1
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \ 2
  --to=<path_to_directory_for_credentials_requests> \ 3
```

- 1 The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2 Specify the location of the **install-config.yaml** file.
- 3 Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.



NOTE

This command might take a few moments to run.

3. Use the **ccoctl** tool to process all **CredentialsRequest** objects by running the following command:

```
$ ccoctl aws create-all \
  --name=<name> \ 1
  --region=<aws_region> \ 2
  --credentials-requests-dir=<path_to_credentials_requests_directory> \ 3
  --output-dir=<path_to_ccoctl_output_dir> \ 4
  --create-private-s3-bucket \ 5
```

- 1 Specify the name used to tag any cloud resources that are created for tracking.
- 2 Specify the AWS region in which cloud resources will be created.
- 3 Specify the directory containing the files for the component **CredentialsRequest** objects.
- 4 Optional: Specify the directory in which you want the **ccoctl** utility to create objects. By default, the utility creates objects in the directory in which the commands are run.
- 5 Optional: By default, the **ccoctl** utility stores the OpenID Connect (OIDC) configuration files in a public S3 bucket and uses the S3 URL as the public OIDC endpoint. To store the OIDC configuration in a private S3 bucket that is accessed by the IAM identity provider through a public CloudFront distribution URL instead, use the **--create-private-s3-bucket** parameter.



NOTE

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

Verification

- To verify that the OpenShift Container Platform secrets are created, list the files in the `<path_to_ccoctl_output_dir>/manifests` directory:

```
$ ls <path_to_ccoctl_output_dir>/manifests
```

Example output

```
cluster-authentication-02-config.yaml
openshift-cloud-credential-operator-cloud-credential-operator-iam-ro-creds-credentials.yaml
openshift-cloud-network-config-controller-cloud-credentials-credentials.yaml
openshift-cluster-api-capamanager-bootstrap-credentials-credentials.yaml
openshift-cluster-csi-drivers-ebs-cloud-credentials-credentials.yaml
openshift-image-registry-installer-cloud-credentials-credentials.yaml
openshift-ingress-operator-cloud-credentials-credentials.yaml
openshift-machine-api-aws-cloud-credentials-credentials.yaml
```

You can verify that the IAM roles are created by querying AWS. For more information, refer to AWS documentation on listing IAM roles.

6.3.5.5.2.2. Creating AWS resources individually

You can use the **ccoctl** tool to create AWS resources individually. This option might be useful for an organization that shares the responsibility for creating these resources among different users or departments.

Otherwise, you can use the **ccoctl aws create-all** command to create the AWS resources automatically. For more information, see "Creating AWS resources with a single command".



NOTE

By default, **ccoctl** creates objects in the directory in which the commands are run. To create the objects in a different directory, use the **--output-dir** flag. This procedure uses `<path_to_ccoctl_output_dir>` to refer to this directory.

Some **ccoctl** commands make AWS API calls to create or modify AWS resources. You can use the **--dry-run** flag to avoid making API calls. Using this flag creates JSON files on the local file system instead. You can review and modify the JSON files and then apply them with the AWS CLI tool using the **--cli-input-json** parameters.

Prerequisites

- Extract and prepare the **ccoctl** binary.

Procedure

1. Generate the public and private RSA key files that are used to set up the OpenID Connect provider for the cluster by running the following command:

```
$ ccoctl aws create-key-pair
```

Example output

```

2021/04/13 11:01:02 Generating RSA keypair
2021/04/13 11:01:03 Writing private key to /<path_to_ccoctl_output_dir>/serviceaccount-
signer.private
2021/04/13 11:01:03 Writing public key to /<path_to_ccoctl_output_dir>/serviceaccount-
signer.public
2021/04/13 11:01:03 Copying signing key for use by installer

```

where **serviceaccount-signer.private** and **serviceaccount-signer.public** are the generated key files.

This command also creates a private key that the cluster requires during installation in **/<path_to_ccoctl_output_dir>/tls/bound-service-account-signing-key.key**.

2. Create an OpenID Connect identity provider and S3 bucket on AWS by running the following command:

```

$ ccoctl aws create-identity-provider \
  --name=<name> \ 1
  --region=<aws_region> \ 2
  --public-key-file=<path_to_ccoctl_output_dir>/serviceaccount-signer.public 3

```

1 **<name>** is the name used to tag any cloud resources that are created for tracking.

2 **<aws-region>** is the AWS region in which cloud resources will be created.

3 **<path_to_ccoctl_output_dir>** is the path to the public key file that the **ccoctl aws create-key-pair** command generated.

Example output

```

2021/04/13 11:16:09 Bucket <name>-oidc created
2021/04/13 11:16:10 OpenID Connect discovery document in the S3 bucket <name>-oidc at
.well-known/openid-configuration updated
2021/04/13 11:16:10 Reading public key
2021/04/13 11:16:10 JSON web key set (JWKS) in the S3 bucket <name>-oidc at keys.json
updated
2021/04/13 11:16:18 Identity Provider created with ARN: arn:aws:iam::
<aws_account_id>:oidc-provider/<name>-oidc.s3.<aws_region>.amazonaws.com

```

where **openid-configuration** is a discovery document and **keys.json** is a JSON web key set file.

This command also creates a YAML configuration file in **/<path_to_ccoctl_output_dir>/manifests/cluster-authentication-02-config.yaml**. This file sets the issuer URL field for the service account tokens that the cluster generates, so that the AWS IAM identity provider trusts the tokens.

3. Create IAM roles for each component in the cluster:
 - a. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```

$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')

```

- b. Extract the list of **CredentialsRequest** objects from the OpenShift Container Platform release image:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included \1
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \
  \2
  --to=<path_to_directory_for_credentials_requests> \3
```

- 1 The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2 Specify the location of the **install-config.yaml** file.
- 3 Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

- c. Use the **ccoctl** tool to process all **CredentialsRequest** objects by running the following command:

```
$ ccoctl aws create-iam-roles \
  --name=<name> \
  --region=<aws_region> \
  --credentials-requests-dir=<path_to_credentials_requests_directory> \
  --identity-provider-arn=arn:aws:iam::<aws_account_id>:oidc-provider/<name>-oidc.s3.
  <aws_region>.amazonaws.com
```



NOTE

For AWS environments that use alternative IAM API endpoints, such as GovCloud, you must also specify your region with the **--region** parameter.

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

For each **CredentialsRequest** object, **ccoctl** creates an IAM role with a trust policy that is tied to the specified OIDC identity provider, and a permissions policy as defined in each **CredentialsRequest** object from the OpenShift Container Platform release image.

Verification

- To verify that the OpenShift Container Platform secrets are created, list the files in the **<path_to_ccoctl_output_dir>/manifests** directory:

```
$ ls <path_to_ccoctl_output_dir>/manifests
```

Example output

```
cluster-authentication-02-config.yaml
openshift-cloud-credential-operator-cloud-credential-operator-iam-ro-creds-credentials.yaml
```



```

openshift-cloud-network-config-controller-cloud-credentials-credentials.yaml
openshift-cluster-api-capamanager-bootstrap-credentials-credentials.yaml
openshift-cluster-csi-drivers-ebs-cloud-credentials-credentials.yaml
openshift-image-registry-installer-cloud-credentials-credentials.yaml
openshift-ingress-operator-cloud-credentials-credentials.yaml
openshift-machine-api-aws-cloud-credentials-credentials.yaml

```

You can verify that the IAM roles are created by querying AWS. For more information, refer to AWS documentation on listing IAM roles.

6.3.5.5.2.3. Incorporating the Cloud Credential Operator utility manifests

To implement short-term security credentials managed outside the cluster for individual components, you must move the manifest files that the Cloud Credential Operator utility (**ccoctl**) created to the correct directories for the installation program.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have configured the Cloud Credential Operator utility (**ccoctl**).
- You have created the cloud provider resources that are required for your cluster with the **ccoctl** utility.

Procedure

1. If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```

apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...

```

2. If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

3. Copy the manifests that the **ccoctl** utility generated to the **manifests** directory that the installation program created by running the following command:


```
$ cp /<path_to_ccoctl_output_dir>/manifests/* ./manifests/
```

4. Copy the private key that the **ccoctl** utility generated in the **tls** directory to the installation directory by running the following command:

```
$ cp -a /<path_to_ccoctl_output_dir>/tls .
```

6.3.5.6. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.

Procedure

1. Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

- 1** For **<installation_directory>**, specify the location of your customized **./install-config.yaml** file.
- 2** To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

2. Optional: Remove or disable the **AdministratorAccess** policy from the IAM account that you used to install the cluster.




NOTE

The elevated permissions provided by the **AdministratorAccess** policy are required only during installation.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to **<installation_directory>/openshift_install.log**.



IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

6.3.5.7. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

6.3.5.8. Disabling the default OperatorHub catalog sources

Operator catalogs that source content provided by Red Hat and community projects are configured for OperatorHub by default during an OpenShift Container Platform installation. In a restricted network environment, you must disable the default catalogs as a cluster administrator.

Procedure

- Disable the sources for the default catalogs by adding **disableAllDefaultSources: true** to the **OperatorHub** object:

```
$ oc patch OperatorHub cluster --type json \
-p '[{"op": "add", "path": "/spec/disableAllDefaultSources", "value": true}]'
```

TIP

Alternatively, you can use the web console to manage catalog sources. From the **Administration** → **Cluster Settings** → **Configuration** → **OperatorHub** page, click the **Sources** tab, where you can create, update, delete, disable, and enable individual sources.

6.3.5.9. Next steps

- [Validate an installation](#).
- [Customize your cluster](#).
- [Configure image streams](#) for the Cluster Samples Operator and the **must-gather** tool.
- Learn how to [use Operator Lifecycle Manager \(OLM\) on restricted networks](#).
- If the mirror registry that you used to install your cluster has a trusted CA, add it to the cluster by [configuring additional trust stores](#).
- If necessary, you can [opt out of remote health reporting](#).

6.3.6. Installing a cluster on AWS into an existing VPC

In OpenShift Container Platform version 4.16, you can install a cluster into an existing Amazon Virtual Private Cloud (VPC) on Amazon Web Services (AWS). The installation program provisions the rest of the required infrastructure, which you can further customize. To customize the installation, you modify parameters in the **install-config.yaml** file before you install the cluster.

6.3.6.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).

- You [configured an AWS account](#) to host the cluster.
- If the existing VPC is owned by a different account than the cluster, you [shared the VPC](#) between accounts.



IMPORTANT

If you have an AWS profile stored on your computer, it must not use a temporary session token that you generated while using a multi-factor authentication device. The cluster continues to use your current AWS credentials to create AWS resources for the entire life of the cluster, so you must use long-term credentials. To generate appropriate keys, see [Managing Access Keys for IAM Users](#) in the AWS documentation. You can supply the keys when you run the installation program.

- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.

6.3.6.2. About using a custom VPC

In OpenShift Container Platform 4.16, you can deploy a cluster into existing subnets in an existing Amazon Virtual Private Cloud (VPC) in Amazon Web Services (AWS). By deploying OpenShift Container Platform into an existing AWS VPC, you might be able to avoid limit constraints in new accounts or more easily abide by the operational constraints that your company's guidelines set. If you cannot obtain the infrastructure creation permissions that are required to create the VPC yourself, use this installation option.

Because the installation program cannot know what other components are also in your existing subnets, it cannot choose subnet CIDRs and so forth on your behalf. You must configure networking for the subnets that you install your cluster to yourself.

6.3.6.2.1. Requirements for using your VPC

The installation program no longer creates the following components:

- Internet gateways
- NAT gateways
- Subnets
- Route tables
- VPCs
- VPC DHCP options
- VPC endpoints



NOTE

The installation program requires that you use the cloud-provided DNS server. Using a custom DNS server is not supported and causes the installation to fail.

If you use a custom VPC, you must correctly configure it and its subnets for the installation program and the cluster to use. See [Amazon VPC console wizard configurations](#) and [Work with VPCs and subnets](#) in the AWS documentation for more information on creating and managing an AWS VPC.

The installation program cannot:

- Subdivide network ranges for the cluster to use.
- Set route tables for the subnets.
- Set VPC options like DHCP.

You must complete these tasks before you install the cluster. See [VPC networking components](#) and [Route tables for your VPC](#) for more information on configuring networking in an AWS VPC.

Your VPC must meet the following characteristics:

- Create a public and private subnet for each availability zone that your cluster uses. Each availability zone can contain no more than one public and one private subnet. For an example of this type of configuration, see [VPC with public and private subnets \(NAT\)](#) in the AWS documentation.
Record each subnet ID. Completing the installation requires that you enter these values in the **platform** section of the **install-config.yaml** file. See [Finding a subnet ID](#) in the AWS documentation.
- The VPC's CIDR block must contain the **Networking.MachineCIDR** range, which is the IP address pool for cluster machines. The subnet CIDR blocks must belong to the machine CIDR that you specify.
- The VPC must have a public internet gateway attached to it. For each availability zone:
 - The public subnet requires a route to the internet gateway.
 - The public subnet requires a NAT gateway with an EIP address.
 - The private subnet requires a route to the NAT gateway in public subnet.
- The VPC must not use the **kubernetes.io/cluster/.*: owned, Name**, and **openshift.io/cluster** tags.
The installation program modifies your subnets to add the **kubernetes.io/cluster/.*: shared** tag, so your subnets must have at least one free tag slot available for it. See [Tag Restrictions](#) in the AWS documentation to confirm that the installation program can add a tag to each subnet that you specify. You cannot use a **Name** tag, because it overlaps with the EC2 **Name** field and the installation fails.
- If you want to extend your OpenShift Container Platform cluster into an AWS Outpost and have an existing Outpost subnet, the existing subnet must use the **kubernetes.io/cluster/unmanaged: true** tag. If you do not apply this tag, the installation might fail due to the Cloud Controller Manager creating a service load balancer in the Outpost subnet, which is an unsupported configuration.
- You must enable the **enableDnsSupport** and **enableDnsHostnames** attributes in your VPC, so that the cluster can use the Route 53 zones that are attached to the VPC to resolve cluster's internal DNS records. See [DNS Support in Your VPC](#) in the AWS documentation.
If you prefer to use your own Route 53 hosted private zone, you must associate the existing hosted zone with your VPC prior to installing a cluster. You can define your hosted zone using the **platform.aws.hostedZone** and **platform.aws.hostedZoneRole** fields in the **install-**

config.yaml file. You can use a private hosted zone from another account by sharing it with the account where you install the cluster. If you use a private hosted zone from another account, you must use the **Passthrough** or **Manual** credentials mode.

If you are working in a disconnected environment, you are unable to reach the public IP addresses for EC2, ELB, and S3 endpoints. Depending on the level to which you want to restrict internet traffic during the installation, the following configuration options are available:

Option 1: Create VPC endpoints

Create a VPC endpoint and attach it to the subnets that the clusters are using. Name the endpoints as follows:

- **ec2.<aws_region>.amazonaws.com**
- **elasticloadbalancing.<aws_region>.amazonaws.com**
- **s3.<aws_region>.amazonaws.com**

With this option, network traffic remains private between your VPC and the required AWS services.

Option 2: Create a proxy without VPC endpoints

As part of the installation process, you can configure an HTTP or HTTPS proxy. With this option, internet traffic goes through the proxy to reach the required AWS services.

Option 3: Create a proxy with VPC endpoints

As part of the installation process, you can configure an HTTP or HTTPS proxy with VPC endpoints. Create a VPC endpoint and attach it to the subnets that the clusters are using. Name the endpoints as follows:

- **ec2.<aws_region>.amazonaws.com**
- **elasticloadbalancing.<aws_region>.amazonaws.com**
- **s3.<aws_region>.amazonaws.com**

When configuring the proxy in the **install-config.yaml** file, add these endpoints to the **noProxy** field. With this option, the proxy prevents the cluster from accessing the internet directly. However, network traffic remains private between your VPC and the required AWS services.

Required VPC components

You must provide a suitable VPC and subnets that allow communication to your machines.

Component	AWS type	Description
VPC	<ul style="list-style-type: none"> • AWS::EC2::VPC • AWS::EC2::VPCEndpoint 	You must provide a public VPC for the cluster to use. The VPC uses an endpoint that references the route tables for each subnet to improve communication with the registry that is hosted in S3.

Component	AWS type	Description	
Public subnets	<ul style="list-style-type: none"> ● AWS::EC2::Subnet ● AWS::EC2::SubnetNetworkACLAssociation 	Your VPC must have public subnets for between 1 and 3 availability zones and associate them with appropriate Ingress rules.	
Internet gateway	<ul style="list-style-type: none"> ● AWS::EC2::InternetGateway ● AWS::EC2::VPCGatewayAttachment ● AWS::EC2::RouteTable ● AWS::EC2::Route ● AWS::EC2::SubnetRouteTableAssociation ● AWS::EC2::NatGateway ● AWS::EC2::EIP 	You must have a public internet gateway, with public routes, attached to the VPC. In the provided templates, each public subnet has a NAT gateway with an EIP address. These NAT gateways allow cluster resources, like private subnet instances, to reach the internet and are not required for some restricted network or proxy scenarios.	
Network access control	<ul style="list-style-type: none"> ● AWS::EC2::NetworkACL ● AWS::EC2::NetworkACLEntry 	You must allow the VPC to access the following ports:	
		Port	Reason
		80	Inbound HTTP traffic
		443	Inbound HTTPS traffic
		22	Inbound SSH traffic
		1024 - 65535	Inbound ephemeral traffic
	0 - 65535	Outbound ephemeral traffic	
Private subnets	<ul style="list-style-type: none"> ● AWS::EC2::Subnet ● AWS::EC2::RouteTable ● AWS::EC2::SubnetRouteTableAssociation 	Your VPC can have private subnets. The provided CloudFormation templates can create private subnets for between 1 and 3 availability zones. If you use private subnets, you must provide appropriate routes and tables for them.	

6.3.6.2.2. VPC validation

To ensure that the subnets that you provide are suitable, the installation program confirms the following data:

- All the subnets that you specify exist.
- You provide private subnets.
- The subnet CIDRs belong to the machine CIDR that you specified.
- You provide subnets for each availability zone. Each availability zone contains no more than one public and one private subnet. If you use a private cluster, provide only a private subnet for each availability zone. Otherwise, provide exactly one public and private subnet for each availability zone.
- You provide a public subnet for each private subnet availability zone. Machines are not provisioned in availability zones that you do not provide private subnets for.

If you destroy a cluster that uses an existing VPC, the VPC is not deleted. When you remove the OpenShift Container Platform cluster from a VPC, the **kubernetes.io/cluster/.*: shared** tag is removed from the subnets that it used.

6.3.6.2.3. Division of permissions

Starting with OpenShift Container Platform 4.3, you do not need all of the permissions that are required for an installation program-provisioned infrastructure cluster to deploy a cluster. This change mimics the division of permissions that you might have at your company: some individuals can create different resource in your clouds than others. For example, you might be able to create application-specific items, like instances, buckets, and load balancers, but not networking-related components such as VPCs, subnets, or ingress rules.

The AWS credentials that you use when you create your cluster do not need the networking permissions that are required to make VPCs and core networking components within the VPC, such as subnets, routing tables, internet gateways, NAT, and VPN. You still need permission to make the application resources that the machines within the cluster require, such as ELBs, security groups, S3 buckets, and nodes.

6.3.6.2.4. Isolation between clusters

If you deploy OpenShift Container Platform to an existing network, the isolation of cluster services is reduced in the following ways:

- You can install multiple OpenShift Container Platform clusters in the same VPC.
- ICMP ingress is allowed from the entire network.
- TCP 22 ingress (SSH) is allowed to the entire network.
- Control plane TCP 6443 ingress (Kubernetes API) is allowed to the entire network.
- Control plane TCP 22623 ingress (MCS) is allowed to the entire network.

6.3.6.2.5. Optional: AWS security groups

By default, the installation program creates and attaches security groups to control plane and compute machines. The rules associated with the default security groups cannot be modified.

However, you can apply additional existing AWS security groups, which are associated with your existing VPC, to control plane and compute machines. Applying custom security groups can help you meet the security needs of your organization, in such cases where you need to control the incoming or outgoing traffic of these machines.

As part of the installation process, you apply custom security groups by modifying the **install-config.yaml** file before deploying the cluster.

For more information, see "Applying existing AWS security groups to the cluster".

6.3.6.2.6. Modifying trust policy when installing into a shared VPC

If you install your cluster using a shared VPC, you can use the **Passthrough** or **Manual** credentials mode. You must add the IAM role used to install the cluster as a principal in the trust policy of the account that owns the VPC.

If you use **Passthrough** mode, add the Amazon Resource Name (ARN) of the account that creates the cluster, such as **arn:aws:iam::123456789012:user/clustercreator**, to the trust policy as a principal.

If you use **Manual** mode, add the ARN of the account that creates the cluster as well as the ARN of the ingress operator role in the cluster owner account, such as **arn:aws:iam::123456789012:role/<cluster-name>-openshift-ingress-operator-cloud-credentials**, to the trust policy as principals.

You must add the following actions to the policy:

Example 6.28. Required actions for shared VPC installation

- **route53:ChangeResourceRecordSets**
- **route53:ListHostedZones**
- **route53:ListHostedZonesByName**
- **route53:ListResourceRecordSets**
- **route53:ChangeTagsForResource**
- **route53:GetAccountLimit**
- **route53:GetChange**
- **route53:GetHostedZone**
- **route53:ListTagsForResource**
- **route53:UpdateHostedZoneComment**
- **tag:GetResources**
- **tag:UntagResources**

6.3.6.3. Creating the installation configuration file

You can customize the OpenShift Container Platform cluster you install on Amazon Web Services (AWS).

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create the **install-config.yaml** file.
 - a. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install create install-config --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the directory name to store the files that the installation program creates.

When specifying the directory:

- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
- Use an empty directory. Some installation assets, such as bootstrap X.509 certificates, have short expiration intervals, therefore you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

- b. At the prompts, provide the configuration details for your cloud:
 - i. Optional: Select an SSH key to use to access your cluster machines.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- ii. Select **AWS** as the platform to target.
 - iii. If you do not have an Amazon Web Services (AWS) profile stored on your computer, enter the AWS access key ID and secret access key for the user that you configured to run the installation program.
 - iv. Select the AWS region to deploy the cluster to.
 - v. Select the base domain for the Route 53 service that you configured for your cluster.
 - vi. Enter a descriptive name for your cluster.
2. Modify the **install-config.yaml** file. You can find more information about the available parameters in the "Installation configuration parameters" section.
 3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

Additional resources

- [Installation configuration parameters for AWS](#)

6.3.6.3.1. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 6.15. Minimum resource requirements

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Bootstrap	RHCOS	4	16 GB	100 GB	300
Control plane	RHCOS	4	16 GB	100 GB	300
Compute	RHCOS, RHEL 8.6 and later [3]	2	8 GB	100 GB	300

1. One vCPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio: $(\text{threads per core} \times \text{cores}) \times \text{sockets} = \text{vCPUs}$.
2. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes which require a 10 ms p99 fsync duration. Note that on many cloud platforms, storage size and IOPS scale together, so you might need to over-allocate storage volume to obtain sufficient performance.
3. As with all user-provisioned installations, if you choose to use RHEL compute machines in your cluster, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Use of RHEL 7 compute machines is deprecated and has been removed in OpenShift Container Platform 4.10 and later.

**NOTE**

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

6.3.6.3.2. Tested instance types for AWS

The following Amazon Web Services (AWS) instance types have been tested with OpenShift Container Platform.

**NOTE**

Use the machine types included in the following charts for your AWS instances. If you use an instance type that is not listed in the chart, ensure that the instance size you use matches the minimum resource requirements that are listed in the section named "Minimum resource requirements for cluster installation".

Example 6.29. Machine types based on 64-bit x86 architecture

- **c4.***
- **c5.***
- **c5a.***
- **i3.***
- **m4.***
- **m5.***
- **m5a.***
- **m6a.***
- **m6i.***

- **r4.***
- **r5.***
- **r5a.***
- **r6i.***
- **t3.***
- **t3a.***

6.3.6.3.3. Tested instance types for AWS on 64-bit ARM infrastructures

The following Amazon Web Services (AWS) 64-bit ARM instance types have been tested with OpenShift Container Platform.



NOTE


Use the machine types included in the following charts for your AWS ARM instances. If you use an instance type that is not listed in the chart, ensure that the instance size you use matches the minimum resource requirements that are listed in "Minimum resource requirements for cluster installation".

Example 6.30. Machine types based on 64-bit ARM architecture

- **c6g.***
- **m6g.***

6.3.6.3.4. Sample customized install-config.yaml file for AWS

You can customize the installation configuration file (**install-config.yaml**) to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.



IMPORTANT

This sample YAML file is provided for reference only. You must obtain your **install-config.yaml** file by using the installation program and modify it.

```
apiVersion: v1
baseDomain: example.com 1
credentialsMode: Mint 2
controlPlane: 3 4
  hyperthreading: Enabled 5
name: master
platform:
  aws:
    zones:
      - us-west-2a
      - us-west-2b
```

```

rootVolume:
  iops: 4000
  size: 500
  type: io1 6
metadataService:
  authentication: Optional 7
  type: m6i.xlarge
replicas: 3
compute: 8
- hyperthreading: Enabled 9
  name: worker
  platform:
    aws:
      rootVolume:
        iops: 2000
        size: 500
        type: io1 10
      metadataService:
        authentication: Optional 11
        type: c5.4xlarge
      zones:
        - us-west-2c
      replicas: 3
  metadata:
    name: test-cluster 12
  networking:
    clusterNetwork:
      - cidr: 10.128.0.0/14
      hostPrefix: 23
    machineNetwork:
      - cidr: 10.0.0.0/16
    networkType: OVNKubernetes 13
    serviceNetwork:
      - 172.30.0.0/16
  platform:
    aws:
      region: us-west-2 14
      propagateUserTags: true 15
      userTags:
        adminContact: jdoe
        costCenter: 7536
      subnets: 16
      - subnet-1
      - subnet-2
      - subnet-3
      amiID: ami-0c5d3e03c0ab9b19a 17
      serviceEndpoints: 18
      - name: ec2
        url: https://vpce-id.ec2.us-west-2.vpce.amazonaws.com
      hostedZone: Z3URY6TWQ91KVV 19
  fips: false 20
  sshKey: ssh-ed25519 AAAA... 21
  pullSecret: '{"auths": ...}' 22

```

- 1 12 14 22 Required. The installation program prompts you for this value.
- 2 Optional: Add this parameter to force the Cloud Credential Operator (CCO) to use the specified mode. By default, the CCO uses the root credentials in the **kube-system** namespace to dynamically try to determine the capabilities of the credentials. For details about CCO modes, see the "About the Cloud Credential Operator" section in the *Authentication and authorization* guide.
- 3 8 15 If you do not provide these parameters and values, the installation program provides the default value.
- 4 The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.
- 5 9 Whether to enable or disable simultaneous multithreading, or **hyperthreading**. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores. You can disable it by setting the parameter value to **Disabled**. If you disable simultaneous multithreading in some cluster machines, you must disable it in all cluster machines.



IMPORTANT

If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance. Use larger instance types, such as **m4.2xlarge** or **m5.2xlarge**, for your machines if you disable simultaneous multithreading.

- 6 10 To configure faster storage for etcd, especially for larger clusters, set the storage type as **io1** and set **iopts** to **2000**.
- 7 11 Whether to require the [Amazon EC2 Instance Metadata Service v2](#) (IMDSv2). To require IMDSv2, set the parameter value to **Required**. To allow the use of both IMDSv1 and IMDSv2, set the parameter value to **Optional**. If no value is specified, both IMDSv1 and IMDSv2 are allowed.



NOTE

The IMDS configuration for control plane machines that is set during cluster installation can only be changed by using the AWS CLI. The IMDS configuration for compute machines can be changed by using compute machine sets.

- 13 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 16 If you provide your own VPC, specify subnets for each availability zone that your cluster uses.
- 17 The ID of the AMI used to boot machines for the cluster. If set, the AMI must belong to the same region as the cluster.
- 18 The AWS service endpoints. Custom endpoints are required when installing to an unknown AWS region. The endpoint URL must use the **https** protocol and the host must trust the certificate.
- 19 The ID of your existing Route 53 private hosted zone. Providing an existing hosted zone requires that you supply your own VPC and the hosted zone is already associated with the VPC prior to installing your cluster. If undefined, the installation program creates a new hosted zone.

- 20 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container




IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 21 You can optionally provide the **sshKey** value that you use to access the machines in your cluster.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

6.3.6.3.5. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure


1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```

apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: ec2.<aws_region>.amazonaws.com,elasticloadbalancing.
<aws_region>.amazonaws.com,s3.<aws_region>.amazonaws.com 3
additionalTrustBundle: | 4
  -----BEGIN CERTIFICATE-----
  <MY_TRUSTED_CA_CERT>
  -----END CERTIFICATE-----
additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5

```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations. If you have added the Amazon **EC2,Elastic Load Balancing**, and **S3** VPC endpoints to your VPC, you must add these endpoints to the **noProxy** field.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.



NOTE

The installation program does not support the proxy **readinessEndpoints** field.



NOTE

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.



NOTE

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

6.3.6.3.6. Applying existing AWS security groups to the cluster

Applying existing AWS security groups to your control plane and compute machines can help you meet the security needs of your organization, in such cases where you need to control the incoming or outgoing traffic of these machines.

Prerequisites

- You have created the security groups in AWS. For more information, see the AWS documentation about working with [security groups](#).
- The security groups must be associated with the existing VPC that you are deploying the cluster to. The security groups cannot be associated with another VPC.
- You have an existing **install-config.yaml** file.

Procedure

1. In the **install-config.yaml** file, edit the **compute.platform.aws.additionalSecurityGroupIDs** parameter to specify one or more custom security groups for your compute machines.
2. Edit the **controlPlane.platform.aws.additionalSecurityGroupIDs** parameter to specify one or more custom security groups for your control plane machines.
3. Save the file and reference it when deploying the cluster.

Sample **install-config.yaml** file that specifies custom security groups

```
# ...
compute:
- hyperthreading: Enabled
  name: worker
  platform:
    aws:
      additionalSecurityGroupIDs:
        - sg-1 1
        - sg-2
  replicas: 3
controlPlane:
  hyperthreading: Enabled
  name: master
  platform:
    aws:
      additionalSecurityGroupIDs:
        - sg-3
        - sg-4
```

```

replicas: 3
platform:
aws:
  region: us-east-1
  subnets: 2
    - subnet-1
    - subnet-2
    - subnet-3

```

- 1 Specify the name of the security group as it appears in the Amazon EC2 console, including the **sg** prefix.
- 2 Specify subnets for each availability zone that your cluster uses.

6.3.6.4. Alternatives to storing administrator-level secrets in the kube-system project

By default, administrator secrets are stored in the **kube-system** project. If you configured the **credentialsMode** parameter in the **install-config.yaml** file to **Manual**, you must use one of the following alternatives:

- To manage long-term cloud credentials manually, follow the procedure in [Manually creating long-term credentials](#).
- To implement short-term credentials that are managed outside the cluster for individual components, follow the procedures in [Configuring an AWS cluster to use short-term credentials](#).

6.3.6.4.1. Manually creating long-term credentials

The Cloud Credential Operator (CCO) can be put into manual mode prior to installation in environments where the cloud identity and access management (IAM) APIs are not reachable, or the administrator prefers not to store an administrator-level credential secret in the cluster **kube-system** namespace.

Procedure

1. If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```

apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...

```

2. If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

- Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

- Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included \ ❶
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \ ❷
  --to=<path_to_directory_for_credentials_requests> \ ❸
```

- The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- Specify the location of the **install-config.yaml** file.
- Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

This command creates a YAML file for each **CredentialsRequest** object.

Sample **CredentialsRequest** object

```
apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: AWSProviderSpec
    statementEntries:
      - effect: Allow
        action:
          - iam:GetUser
          - iam:GetUserPolicy
          - iam:ListAccessKeys
        resource: "*"
  ...
```

- Create YAML files for secrets in the **openshift-install** manifests directory that you generated previously. The secrets must be stored using the namespace and secret name defined in the **spec.secretRef** for each **CredentialsRequest** object.

Sample **CredentialsRequest** object with secrets

```
apiVersion: cloudcredential.openshift.io/v1
```

```

kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: AWSProviderSpec
    statementEntries:
      - effect: Allow
        action:
          - s3:CreateBucket
          - s3>DeleteBucket
        resource: "*"
    ...
  secretRef:
    name: <component_secret>
    namespace: <component_namespace>
  ...


```

Sample Secret object

```

apiVersion: v1
kind: Secret
metadata:
  name: <component_secret>
  namespace: <component_namespace>
data:
  aws_access_key_id: <base64_encoded_aws_access_key_id>
  aws_secret_access_key: <base64_encoded_aws_secret_access_key>

```



IMPORTANT

Before upgrading a cluster that uses manually maintained credentials, you must ensure that the CCO is in an upgradeable state.

6.3.6.4.2. Configuring an AWS cluster to use short-term credentials

To install a cluster that is configured to use the AWS Security Token Service (STS), you must configure the CCO utility and create the required AWS resources for your cluster.

6.3.6.4.2.1. Configuring the Cloud Credential Operator utility

To create and manage cloud credentials from outside of the cluster when the Cloud Credential Operator (CCO) is operating in manual mode, extract and prepare the CCO utility (**ccoctl**) binary.



NOTE

The **ccoctl** utility is a Linux binary that must run in a Linux environment.

Prerequisites

- You have access to an OpenShift Container Platform account with cluster administrator access.

- You have installed the OpenShift CLI (**oc**).
- You have created an AWS account for the **ccoctl** utility to use with the following permissions:

Example 6.31. Required AWS permissions

Required iam permissions

- **iam:CreateOpenIDConnectProvider**
- **iam:CreateRole**
- **iam>DeleteOpenIDConnectProvider**
- **iam>DeleteRole**
- **iam>DeleteRolePolicy**
- **iam:GetOpenIDConnectProvider**
- **iam:GetRole**
- **iam:GetUser**
- **iam:ListOpenIDConnectProviders**
- **iam:ListRolePolicies**
- **iam:ListRoles**
- **iam:PutRolePolicy**
- **iam:TagOpenIDConnectProvider**
- **iam:TagRole**

Required s3 permissions

- **s3:CreateBucket**
- **s3>DeleteBucket**
- **s3>DeleteObject**
- **s3:GetBucketAcl**
- **s3:GetBucketTagging**
- **s3:GetObject**
- **s3:GetObjectAcl**
- **s3:GetObjectTagging**
- **s3:ListBucket**
- **s3:PutBucketAcl**
- **s3:PutBucketPolicy**

- **s3:PutBucketPublicAccessBlock**
- **s3:PutBucketTagging**
- **s3:PutObject**
- **s3:PutObjectAcl**
- **s3:PutObjectTagging**

Required cloudfront permissions

- **cloudfront:ListCloudFrontOriginAccessIdentities**
- **cloudfront:ListDistributions**
- **cloudfront:ListTagsForResource**

If you plan to store the OIDC configuration in a private S3 bucket that is accessed by the IAM identity provider through a public CloudFront distribution URL, the AWS account that runs the **ccoctl** utility requires the following additional permissions:

Example 6.32. Additional permissions for a private S3 bucket with CloudFront

- **cloudfront:CreateCloudFrontOriginAccessIdentity**
- **cloudfront:CreateDistribution**
- **cloudfront>DeleteCloudFrontOriginAccessIdentity**
- **cloudfront>DeleteDistribution**
- **cloudfront:GetCloudFrontOriginAccessIdentity**
- **cloudfront:GetCloudFrontOriginAccessIdentityConfig**
- **cloudfront:GetDistribution**
- **cloudfront:TagResource**
- **cloudfront:UpdateDistribution**



NOTE

These additional permissions support the use of the **--create-private-s3-bucket** option when processing credentials requests with the **ccoctl aws create-all** command.


Procedure

1. Set a variable for the OpenShift Container Platform release image by running the following command:


```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

- Obtain the CCO container image from the OpenShift Container Platform release image by running the following command:

```
$ CCO_IMAGE=$(oc adm release info --image-for='cloud-credential-operator'
$RELEASE_IMAGE -a ~/.pull-secret)
```



NOTE

Ensure that the architecture of the **\$RELEASE_IMAGE** matches the architecture of the environment in which you will use the **ccoctl** tool.

- Extract the **ccoctl** binary from the CCO container image within the OpenShift Container Platform release image by running the following command:

```
$ oc image extract $CCO_IMAGE \
--file="/usr/bin/ccoctl.<rhel_version>" \
-a ~/.pull-secret
```

- For **<rhel_version>**, specify the value that corresponds to the version of Red Hat Enterprise Linux (RHEL) that the host uses. If no value is specified, **ccoctl.rhel8** is used by default. The following values are valid:

- rhel8**: Specify this value for hosts that use RHEL 8.
- rhel9**: Specify this value for hosts that use RHEL 9.

- Change the permissions to make **ccoctl** executable by running the following command:

```
$ chmod 775 ccoctl.<rhel_version>
```

Verification

- To verify that **ccoctl** is ready to use, display the help file by running the following command:

```
$ ccoctl --help
```

Output of **ccoctl --help**

```
OpenShift credentials provisioning tool
```

```
Usage:
```

```
ccoctl [command]
```

```
Available Commands:
```

```
aws      Manage credentials objects for AWS cloud
azure    Manage credentials objects for Azure
gcp      Manage credentials objects for Google cloud
help     Help about any command
ibmcloud Manage credentials objects for IBM Cloud
nutanix  Manage credentials objects for Nutanix
```

Flags:

-h, --help help for ccoctl

Use "ccoctl [command] --help" for more information about a command.

6.3.6.4.2.2. Creating AWS resources with the Cloud Credential Operator utility

You have the following options when creating AWS resources:

- You can use the **ccoctl aws create-all** command to create the AWS resources automatically. This is the quickest way to create the resources. See [Creating AWS resources with a single command](#).
- If you need to review the JSON files that the **ccoctl** tool creates before modifying AWS resources, or if the process the **ccoctl** tool uses to create AWS resources automatically does not meet the requirements of your organization, you can create the AWS resources individually. See [Creating AWS resources individually](#).

6.3.6.4.2.2.1. Creating AWS resources with a single command

If the process the **ccoctl** tool uses to create AWS resources automatically meets the requirements of your organization, you can use the **ccoctl aws create-all** command to automate the creation of AWS resources.

Otherwise, you can create the AWS resources individually. For more information, see "Creating AWS resources individually".



NOTE

By default, **ccoctl** creates objects in the directory in which the commands are run. To create the objects in a different directory, use the **--output-dir** flag. This procedure uses **<path_to_ccoctl_output_dir>** to refer to this directory.

Prerequisites

You must have:

- Extracted and prepared the **ccoctl** binary.

Procedure

1. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

2. Extract the list of **CredentialsRequest** objects from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
```

```
--included \ 1
--install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \ 2
--to=<path_to_directory_for_credentials_requests> 3
```

- 1 The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2 Specify the location of the **install-config.yaml** file.
- 3 Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.



NOTE

This command might take a few moments to run.

3. Use the **ccoctl** tool to process all **CredentialsRequest** objects by running the following command:

```
$ ccoctl aws create-all \
  --name=<name> \ 1
  --region=<aws_region> \ 2
  --credentials-requests-dir=<path_to_credentials_requests_directory> \ 3
  --output-dir=<path_to_ccoctl_output_dir> \ 4
  --create-private-s3-bucket 5
```

- 1 Specify the name used to tag any cloud resources that are created for tracking.
- 2 Specify the AWS region in which cloud resources will be created.
- 3 Specify the directory containing the files for the component **CredentialsRequest** objects.
- 4 Optional: Specify the directory in which you want the **ccoctl** utility to create objects. By default, the utility creates objects in the directory in which the commands are run.
- 5 Optional: By default, the **ccoctl** utility stores the OpenID Connect (OIDC) configuration files in a public S3 bucket and uses the S3 URL as the public OIDC endpoint. To store the OIDC configuration in a private S3 bucket that is accessed by the IAM identity provider through a public CloudFront distribution URL instead, use the **--create-private-s3-bucket** parameter.



NOTE

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

Verification

- To verify that the OpenShift Container Platform secrets are created, list the files in the **<path_to_ccoctl_output_dir>/manifests** directory:

```
$ ls <path_to_ccoctl_output_dir>/manifests
```

Example output

```
cluster-authentication-02-config.yaml
openshift-cloud-credential-operator-cloud-credential-operator-iam-ro-creds-credentials.yaml
openshift-cloud-network-config-controller-cloud-credentials-credentials.yaml
openshift-cluster-api-capamanager-bootstrap-credentials-credentials.yaml
openshift-cluster-csi-drivers-ebs-cloud-credentials-credentials.yaml
openshift-image-registry-installer-cloud-credentials-credentials.yaml
openshift-ingress-operator-cloud-credentials-credentials.yaml
openshift-machine-api-aws-cloud-credentials-credentials.yaml
```

You can verify that the IAM roles are created by querying AWS. For more information, refer to AWS documentation on listing IAM roles.

6.3.6.4.2.2. Creating AWS resources individually

You can use the **ccoctl** tool to create AWS resources individually. This option might be useful for an organization that shares the responsibility for creating these resources among different users or departments.

Otherwise, you can use the **ccoctl aws create-all** command to create the AWS resources automatically. For more information, see "Creating AWS resources with a single command".



NOTE

By default, **ccoctl** creates objects in the directory in which the commands are run. To create the objects in a different directory, use the **--output-dir** flag. This procedure uses **<path_to_ccoctl_output_dir>** to refer to this directory.

Some **ccoctl** commands make AWS API calls to create or modify AWS resources. You can use the **--dry-run** flag to avoid making API calls. Using this flag creates JSON files on the local file system instead. You can review and modify the JSON files and then apply them with the AWS CLI tool using the **--cli-input-json** parameters.

Prerequisites

- Extract and prepare the **ccoctl** binary.

Procedure

1. Generate the public and private RSA key files that are used to set up the OpenID Connect provider for the cluster by running the following command:

```
$ ccoctl aws create-key-pair
```

Example output

```
2021/04/13 11:01:02 Generating RSA keypair
2021/04/13 11:01:03 Writing private key to /<path_to_ccoctl_output_dir>/serviceaccount-signer.private
```

```
2021/04/13 11:01:03 Writing public key to /<path_to_ccoctl_output_dir>/serviceaccount-
signer.public
2021/04/13 11:01:03 Copying signing key for use by installer
```

where **serviceaccount-signer.private** and **serviceaccount-signer.public** are the generated key files.

This command also creates a private key that the cluster requires during installation in **/<path_to_ccoctl_output_dir>/tls/bound-service-account-signing-key.key**.

2. Create an OpenID Connect identity provider and S3 bucket on AWS by running the following command:

```
$ ccoctl aws create-identity-provider \
  --name=<name> \ 1
  --region=<aws_region> \ 2
  --public-key-file=<path_to_ccoctl_output_dir>/serviceaccount-signer.public 3
```

- 1** **<name>** is the name used to tag any cloud resources that are created for tracking.
- 2** **<aws-region>** is the AWS region in which cloud resources will be created.
- 3** **<path_to_ccoctl_output_dir>** is the path to the public key file that the **ccoctl aws create-key-pair** command generated.

Example output

```
2021/04/13 11:16:09 Bucket <name>-oidc created
2021/04/13 11:16:10 OpenID Connect discovery document in the S3 bucket <name>-oidc at
.well-known/openid-configuration updated
2021/04/13 11:16:10 Reading public key
2021/04/13 11:16:10 JSON web key set (JWKS) in the S3 bucket <name>-oidc at keys.json
updated
2021/04/13 11:16:18 Identity Provider created with ARN: arn:aws:iam::
<aws_account_id>:oidc-provider/<name>-oidc.s3.<aws_region>.amazonaws.com
```

where **openid-configuration** is a discovery document and **keys.json** is a JSON web key set file.

This command also creates a YAML configuration file in **/<path_to_ccoctl_output_dir>/manifests/cluster-authentication-02-config.yaml**. This file sets the issuer URL field for the service account tokens that the cluster generates, so that the AWS IAM identity provider trusts the tokens.

3. Create IAM roles for each component in the cluster:
 - a. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

- b. Extract the list of **CredentialsRequest** objects from the OpenShift Container Platform release image:

```
$ oc adm release extract \
```

```
--from=$RELEASE_IMAGE \
--credentials-requests \
--included \ 1
--install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \
2
--to=<path_to_directory_for_credentials_requests> 3
```

- 1 The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2 Specify the location of the **install-config.yaml** file.
- 3 Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

- c. Use the **ccoctl** tool to process all **CredentialsRequest** objects by running the following command:

```
$ ccoctl aws create-iam-roles \
--name=<name> \
--region=<aws_region> \
--credentials-requests-dir=<path_to_credentials_requests_directory> \
--identity-provider-arn=arn:aws:iam::<aws_account_id>:oidc-provider/<name>-oidc.s3.
<aws_region>.amazonaws.com
```



NOTE

For AWS environments that use alternative IAM API endpoints, such as GovCloud, you must also specify your region with the **--region** parameter.

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

For each **CredentialsRequest** object, **ccoctl** creates an IAM role with a trust policy that is tied to the specified OIDC identity provider, and a permissions policy as defined in each **CredentialsRequest** object from the OpenShift Container Platform release image.

Verification

- To verify that the OpenShift Container Platform secrets are created, list the files in the **<path_to_ccoctl_output_dir>/manifests** directory:

```
$ ls <path_to_ccoctl_output_dir>/manifests
```

Example output

```
cluster-authentication-02-config.yaml
openshift-cloud-credential-operator-cloud-credential-operator-iam-ro-creds-credentials.yaml
openshift-cloud-network-config-controller-cloud-credentials-credentials.yaml
openshift-cluster-api-capamanager-bootstrap-credentials-credentials.yaml
openshift-cluster-csi-drivers-efs-cloud-credentials-credentials.yaml
```

```

openshift-image-registry-installer-cloud-credentials-credentials.yaml
openshift-ingress-operator-cloud-credentials-credentials.yaml
openshift-machine-api-aws-cloud-credentials-credentials.yaml

```

You can verify that the IAM roles are created by querying AWS. For more information, refer to AWS documentation on listing IAM roles.

6.3.6.4.2.3. Incorporating the Cloud Credential Operator utility manifests

To implement short-term security credentials managed outside the cluster for individual components, you must move the manifest files that the Cloud Credential Operator utility (**ccoctl**) created to the correct directories for the installation program.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have configured the Cloud Credential Operator utility (**ccoctl**).
- You have created the cloud provider resources that are required for your cluster with the **ccoctl** utility.

Procedure

1. If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```

apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...

```

2. If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

3. Copy the manifests that the **ccoctl** utility generated to the **manifests** directory that the installation program created by running the following command:

```
$ cp /<path_to_ccoctl_output_dir>/manifests/* ./manifests/
```

4. Copy the private key that the **ccoctl** utility generated in the **tls** directory to the installation directory by running the following command:

```
$ cp -a /<path_to_ccoctl_output_dir>/tls .
```

6.3.6.5. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.

Procedure

1. Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

- 1** For **<installation_directory>**, specify the location of your customized **./install-config.yaml** file.
- 2** To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

2. Optional: Remove or disable the **AdministratorAccess** policy from the IAM account that you used to install the cluster.



NOTE

The elevated permissions provided by the **AdministratorAccess** policy are required only during installation.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to **<installation_directory>/openshift_install.log**.




IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

6.3.6.6. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

6.3.6.7. Logging in to the cluster by using the web console

The **kubeadmin** user exists by default after an OpenShift Container Platform installation. You can log in to your cluster as the **kubeadmin** user by using the OpenShift Container Platform web console.

Prerequisites

- You have access to the installation host.
- You completed a cluster installation and all cluster Operators are available.

Procedure

1. Obtain the password for the **kubeadmin** user from the **kubeadmin-password** file on the installation host:

```
$ cat <installation_directory>/auth/kubeadmin-password
```



NOTE

Alternatively, you can obtain the **kubeadmin** password from the **<installation_directory>/openshift_install.log** log file on the installation host.

2. List the OpenShift Container Platform web console route:

```
$ oc get routes -n openshift-console | grep 'console-openshift'
```



NOTE

Alternatively, you can obtain the OpenShift Container Platform route from the **<installation_directory>/openshift_install.log** log file on the installation host.

Example output

```
console    console-openshift-console.apps.<cluster_name>.<base_domain>    console
https reencrypt/Redirect None
```

3. Navigate to the route detailed in the output of the preceding command in a web browser and log in as the **kubeadmin** user.

Additional resources

- See [Accessing the web console](#) for more details about accessing and understanding the OpenShift Container Platform web console.

6.3.6.8. Next steps


- [Validating an installation](#).
- [Customize your cluster](#).
- If necessary, you can [opt out of remote health reporting](#).
- If necessary, you can [remove cloud provider credentials](#).
- After installing a cluster on AWS into an existing VPC, you can [extend the AWS VPC cluster into an AWS Outpost](#).

6.3.7. Installing a private cluster on AWS

In OpenShift Container Platform version 4.16, you can install a private cluster into an existing VPC on Amazon Web Services (AWS). The installation program provisions the rest of the required infrastructure, which you can further customize. To customize the installation, you modify parameters in the `install-config.yaml` file before you install the cluster.

6.3.7.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [configured an AWS account](#) to host the cluster.



IMPORTANT

If you have an AWS profile stored on your computer, it must not use a temporary session token that you generated while using a multi-factor authentication device. The cluster continues to use your current AWS credentials to create AWS resources for the entire life of the cluster, so you must use long-term credentials. To generate appropriate keys, see [Managing Access Keys for IAM Users](#) in the AWS documentation. You can supply the keys when you run the installation program.

- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.

6.3.7.2. Private clusters

You can deploy a private OpenShift Container Platform cluster that does not expose external endpoints. Private clusters are accessible from only an internal network and are not visible to the internet.

By default, OpenShift Container Platform is provisioned to use publicly-accessible DNS and endpoints. A private cluster sets the DNS, Ingress Controller, and API server to private when you deploy your cluster. This means that the cluster resources are only accessible from your internal network and are not visible to the internet.



IMPORTANT

If the cluster has any public subnets, load balancer services created by administrators might be publicly accessible. To ensure cluster security, verify that these services are explicitly annotated as private.

To deploy a private cluster, you must:

- Use existing networking that meets your requirements. Your cluster resources might be shared between other clusters on the network.
- Deploy from a machine that has access to:
 - The API services for the cloud to which you provision.
 - The hosts on the network that you provision.
 - The internet to obtain installation media.

You can use any machine that meets these access requirements and follows your company's guidelines. For example, this machine can be a bastion host on your cloud network or a machine that has access to the network through a VPN.

6.3.7.2.1. Private clusters in AWS

To create a private cluster on Amazon Web Services (AWS), you must provide an existing private VPC and subnets to host the cluster. The installation program must also be able to resolve the DNS records that the cluster requires. The installation program configures the Ingress Operator and API server for access from only the private network.

The cluster still requires access to internet to access the AWS APIs.

The following items are not required or created when you install a private cluster:

- Public subnets
- Public load balancers, which support public ingress
- A public Route 53 zone that matches the **baseDomain** for the cluster

The installation program does use the **baseDomain** that you specify to create a private Route 53 zone and the required records for the cluster. The cluster is configured so that the Operators do not create public records for the cluster and all cluster machines are placed in the private subnets that you specify.

6.3.7.2.1.1. Limitations

The ability to add public functionality to a private cluster is limited.

- You cannot make the Kubernetes API endpoints public after installation without taking additional actions, including creating public subnets in the VPC for each availability zone in use, creating a public load balancer, and configuring the control plane security groups to allow traffic from the internet on 6443 (Kubernetes API port).
- If you use a public Service type load balancer, you must tag a public subnet in each availability zone with **kubernetes.io/cluster/<cluster-infra-id>: shared** so that AWS can use them to create public load balancers.

6.3.7.3. About using a custom VPC

In OpenShift Container Platform 4.16, you can deploy a cluster into existing subnets in an existing Amazon Virtual Private Cloud (VPC) in Amazon Web Services (AWS). By deploying OpenShift Container Platform into an existing AWS VPC, you might be able to avoid limit constraints in new accounts or more easily abide by the operational constraints that your company's guidelines set. If you cannot obtain the infrastructure creation permissions that are required to create the VPC yourself, use this installation option.

Because the installation program cannot know what other components are also in your existing subnets, it cannot choose subnet CIDRs and so forth on your behalf. You must configure networking for the subnets that you install your cluster to yourself.

6.3.7.3.1. Requirements for using your VPC

The installation program no longer creates the following components:

- Internet gateways
- NAT gateways
- Subnets
- Route tables
- VPCs
- VPC DHCP options
- VPC endpoints



NOTE

The installation program requires that you use the cloud-provided DNS server. Using a custom DNS server is not supported and causes the installation to fail.

If you use a custom VPC, you must correctly configure it and its subnets for the installation program and the cluster to use. See [Amazon VPC console wizard configurations](#) and [Work with VPCs and subnets](#) in the AWS documentation for more information on creating and managing an AWS VPC.

The installation program cannot:

- Subdivide network ranges for the cluster to use.
- Set route tables for the subnets.
- Set VPC options like DHCP.

You must complete these tasks before you install the cluster. See [VPC networking components](#) and [Route tables for your VPC](#) for more information on configuring networking in an AWS VPC.

Your VPC must meet the following characteristics:

- The VPC must not use the **kubernetes.io/cluster/.*: owned, Name**, and **openshift.io/cluster** tags.
The installation program modifies your subnets to add the **kubernetes.io/cluster/.*: shared**

tag, so your subnets must have at least one free tag slot available for it. See [Tag Restrictions](#) in the AWS documentation to confirm that the installation program can add a tag to each subnet that you specify. You cannot use a **Name** tag, because it overlaps with the EC2 **Name** field and the installation fails.

- If you want to extend your OpenShift Container Platform cluster into an AWS Outpost and have an existing Outpost subnet, the existing subnet must use the **kubernetes.io/cluster/unmanaged: true** tag. If you do not apply this tag, the installation might fail due to the Cloud Controller Manager creating a service load balancer in the Outpost subnet, which is an unsupported configuration.
- You must enable the **enableDnsSupport** and **enableDnsHostnames** attributes in your VPC, so that the cluster can use the Route 53 zones that are attached to the VPC to resolve cluster's internal DNS records. See [DNS Support in Your VPC](#) in the AWS documentation.
If you prefer to use your own Route 53 hosted private zone, you must associate the existing hosted zone with your VPC prior to installing a cluster. You can define your hosted zone using the **platform.aws.hostedZone** and **platform.aws.hostedZoneRole** fields in the **install-config.yaml** file. You can use a private hosted zone from another account by sharing it with the account where you install the cluster. If you use a private hosted zone from another account, you must use the **Passthrough** or **Manual** credentials mode.

If you are working in a disconnected environment, you are unable to reach the public IP addresses for EC2, ELB, and S3 endpoints. Depending on the level to which you want to restrict internet traffic during the installation, the following configuration options are available:

Option 1: Create VPC endpoints

Create a VPC endpoint and attach it to the subnets that the clusters are using. Name the endpoints as follows:

- **ec2.<aws_region>.amazonaws.com**
- **elasticloadbalancing.<aws_region>.amazonaws.com**
- **s3.<aws_region>.amazonaws.com**

With this option, network traffic remains private between your VPC and the required AWS services.

Option 2: Create a proxy without VPC endpoints

As part of the installation process, you can configure an HTTP or HTTPS proxy. With this option, internet traffic goes through the proxy to reach the required AWS services.

Option 3: Create a proxy with VPC endpoints

As part of the installation process, you can configure an HTTP or HTTPS proxy with VPC endpoints. Create a VPC endpoint and attach it to the subnets that the clusters are using. Name the endpoints as follows:

- **ec2.<aws_region>.amazonaws.com**
- **elasticloadbalancing.<aws_region>.amazonaws.com**
- **s3.<aws_region>.amazonaws.com**

When configuring the proxy in the **install-config.yaml** file, add these endpoints to the **noProxy** field. With this option, the proxy prevents the cluster from accessing the internet directly. However, network traffic remains private between your VPC and the required AWS services.

Required VPC components

You must provide a suitable VPC and subnets that allow communication to your machines.

Component	AWS type	Description										
VPC	<ul style="list-style-type: none"> ● AWS::EC2::VPC ● AWS::EC2::VPCEndpoint 	You must provide a public VPC for the cluster to use. The VPC uses an endpoint that references the route tables for each subnet to improve communication with the registry that is hosted in S3.										
Public subnets	<ul style="list-style-type: none"> ● AWS::EC2::Subnet ● AWS::EC2::SubnetNetworkACLAssociation 	Your VPC must have public subnets for between 1 and 3 availability zones and associate them with appropriate Ingress rules.										
Internet gateway	<ul style="list-style-type: none"> ● AWS::EC2::InternetGateway ● AWS::EC2::VPCGatewayAttachment ● AWS::EC2::RouteTable ● AWS::EC2::Route ● AWS::EC2::SubnetRouteTableAssociation ● AWS::EC2::NatGateway ● AWS::EC2::EIP 	You must have a public internet gateway, with public routes, attached to the VPC. In the provided templates, each public subnet has a NAT gateway with an EIP address. These NAT gateways allow cluster resources, like private subnet instances, to reach the internet and are not required for some restricted network or proxy scenarios.										
Network access control	<ul style="list-style-type: none"> ● AWS::EC2::NetworkACL ● AWS::EC2::NetworkACLEntry 	<p>You must allow the VPC to access the following ports:</p> <table border="1"> <thead> <tr> <th>Port</th> <th>Reason</th> </tr> </thead> <tbody> <tr> <td>80</td> <td>Inbound HTTP traffic</td> </tr> <tr> <td>443</td> <td>Inbound HTTPS traffic</td> </tr> <tr> <td>22</td> <td>Inbound SSH traffic</td> </tr> <tr> <td>1024 - 65535</td> <td>Inbound ephemeral traffic</td> </tr> </tbody> </table>	Port	Reason	80	Inbound HTTP traffic	443	Inbound HTTPS traffic	22	Inbound SSH traffic	1024 - 65535	Inbound ephemeral traffic
Port	Reason											
80	Inbound HTTP traffic											
443	Inbound HTTPS traffic											
22	Inbound SSH traffic											
1024 - 65535	Inbound ephemeral traffic											

Component	AWS type	Description	
		0 - 65535	Outbound ephemeral traffic
Private subnets	<ul style="list-style-type: none"> ● AWS::EC2::Subnet ● AWS::EC2::RouteTable ● AWS::EC2::SubnetRouteTableAssociation 	Your VPC can have private subnets. The provided CloudFormation templates can create private subnets for between 1 and 3 availability zones. If you use private subnets, you must provide appropriate routes and tables for them.	

6.3.7.3.2. VPC validation

To ensure that the subnets that you provide are suitable, the installation program confirms the following data:

- All the subnets that you specify exist.
- You provide private subnets.
- The subnet CIDRs belong to the machine CIDR that you specified.
- You provide subnets for each availability zone. Each availability zone contains no more than one public and one private subnet. If you use a private cluster, provide only a private subnet for each availability zone. Otherwise, provide exactly one public and private subnet for each availability zone.
- You provide a public subnet for each private subnet availability zone. Machines are not provisioned in availability zones that you do not provide private subnets for.

If you destroy a cluster that uses an existing VPC, the VPC is not deleted. When you remove the OpenShift Container Platform cluster from a VPC, the **kubernetes.io/cluster/.*: shared** tag is removed from the subnets that it used.

6.3.7.3.3. Division of permissions

Starting with OpenShift Container Platform 4.3, you do not need all of the permissions that are required for an installation program-provisioned infrastructure cluster to deploy a cluster. This change mimics the division of permissions that you might have at your company: some individuals can create different resource in your clouds than others. For example, you might be able to create application-specific items, like instances, buckets, and load balancers, but not networking-related components such as VPCs, subnets, or ingress rules.

The AWS credentials that you use when you create your cluster do not need the networking permissions that are required to make VPCs and core networking components within the VPC, such as subnets, routing tables, internet gateways, NAT, and VPN. You still need permission to make the application resources that the machines within the cluster require, such as ELBs, security groups, S3 buckets, and nodes.

6.3.7.3.4. Isolation between clusters

If you deploy OpenShift Container Platform to an existing network, the isolation of cluster services is reduced in the following ways:

- You can install multiple OpenShift Container Platform clusters in the same VPC.
- ICMP ingress is allowed from the entire network.
- TCP 22 ingress (SSH) is allowed to the entire network.
- Control plane TCP 6443 ingress (Kubernetes API) is allowed to the entire network.
- Control plane TCP 22623 ingress (MCS) is allowed to the entire network.

6.3.7.3.5. Optional: AWS security groups

By default, the installation program creates and attaches security groups to control plane and compute machines. The rules associated with the default security groups cannot be modified.

However, you can apply additional existing AWS security groups, which are associated with your existing VPC, to control plane and compute machines. Applying custom security groups can help you meet the security needs of your organization, in such cases where you need to control the incoming or outgoing traffic of these machines.

As part of the installation process, you apply custom security groups by modifying the **install-config.yaml** file before deploying the cluster.

For more information, see "Applying existing AWS security groups to the cluster".

6.3.7.4. Manually creating the installation configuration file

Installing the cluster requires that you manually create the installation configuration file.

Prerequisites

- You have an SSH public key on your local machine to provide to the installation program. The key will be used for SSH authentication onto your cluster nodes for debugging and disaster recovery.
- You have obtained the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create an installation directory to store your required installation assets in:

```
$ mkdir <installation_directory>
```



IMPORTANT


You must create a directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

2. Customize the sample **install-config.yaml** file template that is provided and save it in the **<installation_directory>**.

**NOTE**

You must name this configuration file **install-config.yaml**.

3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.

**IMPORTANT**

The **install-config.yaml** file is consumed during the next step of the installation process. You must back it up now.

Additional resources

- [Installation configuration parameters for AWS](#)

6.3.7.4.1. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 6.16. Minimum resource requirements

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Bootstrap	RHCOS	4	16 GB	100 GB	300
Control plane	RHCOS	4	16 GB	100 GB	300
Compute	RHCOS, RHEL 8.6 and later [3]	2	8 GB	100 GB	300

1. One vCPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio: $(\text{threads per core} \times \text{cores}) \times \text{sockets} = \text{vCPUs}$.
2. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes which require a 10 ms p99 fsync duration. Note that on many cloud platforms, storage size and IOPS scale together, so you might need to over-allocate storage volume to obtain sufficient performance.
3. As with all user-provisioned installations, if you choose to use RHEL compute machines in your cluster, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Use of RHEL 7 compute machines is deprecated and has been removed in OpenShift Container Platform 4.10 and later.

**NOTE**

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

6.3.7.4.2. Tested instance types for AWS

The following Amazon Web Services (AWS) instance types have been tested with OpenShift Container Platform.

**NOTE**

Use the machine types included in the following charts for your AWS instances. If you use an instance type that is not listed in the chart, ensure that the instance size you use matches the minimum resource requirements that are listed in the section named "Minimum resource requirements for cluster installation".

Example 6.33. Machine types based on 64-bit x86 architecture

- **c4.***
- **c5.***
- **c5a.***
- **i3.***
- **m4.***
- **m5.***
- **m5a.***
- **m6a.***
- **m6i.***

- r4.*
- r5.*
- r5a.*
- r6i.*
- t3.*
- t3a.*

6.3.7.4.3. Tested instance types for AWS on 64-bit ARM infrastructures

The following Amazon Web Services (AWS) 64-bit ARM instance types have been tested with OpenShift Container Platform.



NOTE


Use the machine types included in the following charts for your AWS ARM instances. If you use an instance type that is not listed in the chart, ensure that the instance size you use matches the minimum resource requirements that are listed in "Minimum resource requirements for cluster installation".

Example 6.34. Machine types based on 64-bit ARM architecture

- c6g.*
- m6g.*

6.3.7.4.4. Sample customized install-config.yaml file for AWS

You can customize the installation configuration file (**install-config.yaml**) to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.



IMPORTANT

This sample YAML file is provided for reference only. You must obtain your **install-config.yaml** file by using the installation program and modify it.

```
apiVersion: v1
baseDomain: example.com 1
credentialsMode: Mint 2
controlPlane: 3 4
  hyperthreading: Enabled 5
name: master
platform:
  aws:
    zones:
      - us-west-2a
      - us-west-2b
```

```

rootVolume:
  iops: 4000
  size: 500
  type: io1 6
metadataService:
  authentication: Optional 7
  type: m6i.xlarge
replicas: 3
compute: 8
- hyperthreading: Enabled 9
name: worker
platform:
  aws:
    rootVolume:
      iops: 2000
      size: 500
      type: io1 10
    metadataService:
      authentication: Optional 11
    type: c5.4xlarge
    zones:
      - us-west-2c
    replicas: 3
metadata:
  name: test-cluster 12
networking:
  clusterNetwork:
    - cidr: 10.128.0.0/14
    hostPrefix: 23
  machineNetwork:
    - cidr: 10.0.0.0/16
  networkType: OVNKubernetes 13
  serviceNetwork:
    - 172.30.0.0/16
platform:
  aws:
    region: us-west-2 14
    propagateUserTags: true 15
    userTags:
      adminContact: jdoe
      costCenter: 7536
    subnets: 16
    - subnet-1
    - subnet-2
    - subnet-3
    amiID: ami-0c5d3e03c0ab9b19a 17
    serviceEndpoints: 18
    - name: ec2
      url: https://vpce-id.ec2.us-west-2.vpce.amazonaws.com
    hostedZone: Z3URY6TWQ91KVV 19
  fips: false 20
  sshKey: ssh-ed25519 AAAA... 21
  publish: Internal 22
  pullSecret: '{"auths": ...}' 23

```

-
- 1 12 14 23 Required. The installation program prompts you for this value.
- 2 Optional: Add this parameter to force the Cloud Credential Operator (CCO) to use the specified mode. By default, the CCO uses the root credentials in the **kube-system** namespace to dynamically try to determine the capabilities of the credentials. For details about CCO modes, see the "About the Cloud Credential Operator" section in the *Authentication and authorization* guide.
- 3 8 15 If you do not provide these parameters and values, the installation program provides the default value.
- 4 The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.
- 5 9 Whether to enable or disable simultaneous multithreading, or **hyperthreading**. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores. You can disable it by setting the parameter value to **Disabled**. If you disable simultaneous multithreading in some cluster machines, you must disable it in all cluster machines.



IMPORTANT

If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance. Use larger instance types, such as **m4.2xlarge** or **m5.2xlarge**, for your machines if you disable simultaneous multithreading.

- 6 10 To configure faster storage for etcd, especially for larger clusters, set the storage type as **io1** and set **iopts** to **2000**.
- 7 11 Whether to require the [Amazon EC2 Instance Metadata Service v2](#) (IMDSv2). To require IMDSv2, set the parameter value to **Required**. To allow the use of both IMDSv1 and IMDSv2, set the parameter value to **Optional**. If no value is specified, both IMDSv1 and IMDSv2 are allowed.



NOTE

The IMDS configuration for control plane machines that is set during cluster installation can only be changed by using the AWS CLI. The IMDS configuration for compute machines can be changed by using compute machine sets.

- 13 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 16 If you provide your own VPC, specify subnets for each availability zone that your cluster uses.
- 17 The ID of the AMI used to boot machines for the cluster. If set, the AMI must belong to the same region as the cluster.
- 18 The AWS service endpoints. Custom endpoints are required when installing to an unknown AWS region. The endpoint URL must use the **https** protocol and the host must trust the certificate.
- 19 The ID of your existing Route 53 private hosted zone. Providing an existing hosted zone requires that you supply your own VPC and the hosted zone is already associated with the VPC prior to installing your cluster. If undefined, the installation program creates a new hosted zone.

- 20 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container



IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 21 You can optionally provide the **sshKey** value that you use to access the machines in your cluster.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- 22 How to publish the user-facing endpoints of your cluster. Set **publish** to **Internal** to deploy a private cluster, which cannot be accessed from the internet. The default value is **External**.

6.3.7.4.5. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure


1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```

apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: ec2.<aws_region>.amazonaws.com,elasticloadbalancing.
<aws_region>.amazonaws.com,s3.<aws_region>.amazonaws.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5

```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations. If you have added the Amazon **EC2,Elastic Load Balancing**, and **S3** VPC endpoints to your VPC, you must add these endpoints to the **noProxy** field.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.



NOTE

The installation program does not support the proxy **readinessEndpoints** field.



NOTE

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```


2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.



NOTE

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

6.3.7.4.6. Applying existing AWS security groups to the cluster

Applying existing AWS security groups to your control plane and compute machines can help you meet the security needs of your organization, in such cases where you need to control the incoming or outgoing traffic of these machines.

Prerequisites

- You have created the security groups in AWS. For more information, see the AWS documentation about working with [security groups](#).
- The security groups must be associated with the existing VPC that you are deploying the cluster to. The security groups cannot be associated with another VPC.
- You have an existing **install-config.yaml** file.

Procedure

1. In the **install-config.yaml** file, edit the **compute.platform.aws.additionalSecurityGroupIDs** parameter to specify one or more custom security groups for your compute machines.
2. Edit the **controlPlane.platform.aws.additionalSecurityGroupIDs** parameter to specify one or more custom security groups for your control plane machines.
3. Save the file and reference it when deploying the cluster.

Sample **install-config.yaml** file that specifies custom security groups

```
# ...
compute:
- hyperthreading: Enabled
  name: worker
  platform:
    aws:
      additionalSecurityGroupIDs:
        - sg-1 1
        - sg-2
  replicas: 3
controlPlane:
  hyperthreading: Enabled
  name: master
  platform:
    aws:
      additionalSecurityGroupIDs:
```

```

- sg-3
- sg-4
replicas: 3
platform:
aws:
  region: us-east-1
  subnets: 2
    - subnet-1
    - subnet-2
    - subnet-3

```

- 1 Specify the name of the security group as it appears in the Amazon EC2 console, including the **sg** prefix.
- 2 Specify subnets for each availability zone that your cluster uses.

6.3.7.5. Alternatives to storing administrator-level secrets in the kube-system project

By default, administrator secrets are stored in the **kube-system** project. If you configured the **credentialsMode** parameter in the **install-config.yaml** file to **Manual**, you must use one of the following alternatives:

- To manage long-term cloud credentials manually, follow the procedure in [Manually creating long-term credentials](#).
- To implement short-term credentials that are managed outside the cluster for individual components, follow the procedures in [Configuring an AWS cluster to use short-term credentials](#).

6.3.7.5.1. Manually creating long-term credentials

The Cloud Credential Operator (CCO) can be put into manual mode prior to installation in environments where the cloud identity and access management (IAM) APIs are not reachable, or the administrator prefers not to store an administrator-level credential secret in the cluster **kube-system** namespace.

Procedure

1. If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```

apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...

```

2. If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

3. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

4. Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included 1 \
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml 2 \
  --to=<path_to_directory_for_credentials_requests> 3
```

- 1** The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2** Specify the location of the **install-config.yaml** file.
- 3** Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

This command creates a YAML file for each **CredentialsRequest** object.

Sample **CredentialsRequest** object

```
apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: AWSPProviderSpec
    statementEntries:
    - effect: Allow
      action:
      - iam:GetUser
      - iam:GetUserPolicy
      - iam:ListAccessKeys
    resource: "*"
  ...
```

5. Create YAML files for secrets in the **openshift-install** manifests directory that you generated previously. The secrets must be stored using the namespace and secret name defined in the **spec.secretRef** for each **CredentialsRequest** object.

Sample **CredentialsRequest** object with secrets

■

```

apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: AWSProviderSpec
    statementEntries:
      - effect: Allow
        action:
          - s3:CreateBucket
          - s3>DeleteBucket
        resource: "*"
    ...
  secretRef:
    name: <component_secret>
    namespace: <component_namespace>
  ...

```

Sample Secret object

```

apiVersion: v1
kind: Secret
metadata:
  name: <component_secret>
  namespace: <component_namespace>
data:
  aws_access_key_id: <base64_encoded_aws_access_key_id>
  aws_secret_access_key: <base64_encoded_aws_secret_access_key>

```



IMPORTANT

Before upgrading a cluster that uses manually maintained credentials, you must ensure that the CCO is in an upgradeable state.

6.3.7.5.2. Configuring an AWS cluster to use short-term credentials

To install a cluster that is configured to use the AWS Security Token Service (STS), you must configure the CCO utility and create the required AWS resources for your cluster.

6.3.7.5.2.1. Configuring the Cloud Credential Operator utility

To create and manage cloud credentials from outside of the cluster when the Cloud Credential Operator (CCO) is operating in manual mode, extract and prepare the CCO utility (**ccoctl**) binary.



NOTE

The **ccoctl** utility is a Linux binary that must run in a Linux environment.

Prerequisites

- You have access to an OpenShift Container Platform account with cluster administrator access.
- You have installed the OpenShift CLI (**oc**).
- You have created an AWS account for the **ccoctl** utility to use with the following permissions:

Example 6.35. Required AWS permissions

Required iam permissions

- **iam:CreateOpenIDConnectProvider**
- **iam:CreateRole**
- **iam>DeleteOpenIDConnectProvider**
- **iam>DeleteRole**
- **iam>DeleteRolePolicy**
- **iam:GetOpenIDConnectProvider**
- **iam:GetRole**
- **iam:GetUser**
- **iam:ListOpenIDConnectProviders**
- **iam:ListRolePolicies**
- **iam:ListRoles**
- **iam:PutRolePolicy**
- **iam:TagOpenIDConnectProvider**
- **iam:TagRole**

Required s3 permissions

- **s3:CreateBucket**
- **s3>DeleteBucket**
- **s3>DeleteObject**
- **s3:GetBucketAcl**
- **s3:GetBucketTagging**
- **s3:GetObject**
- **s3:GetObjectAcl**
- **s3:GetObjectTagging**
- **s3:ListBucket**
- **s3:PutBucketAcl**

- **s3:PutBucketPolicy**
- **s3:PutBucketPublicAccessBlock**
- **s3:PutBucketTagging**
- **s3:PutObject**
- **s3:PutObjectAcl**
- **s3:PutObjectTagging**

Required cloudfront permissions

- **cloudfront:ListCloudFrontOriginAccessIdentities**
- **cloudfront:ListDistributions**
- **cloudfront:ListTagsForResource**

If you plan to store the OIDC configuration in a private S3 bucket that is accessed by the IAM identity provider through a public CloudFront distribution URL, the AWS account that runs the **ccoctl** utility requires the following additional permissions:

Example 6.36. Additional permissions for a private S3 bucket with CloudFront

- **cloudfront:CreateCloudFrontOriginAccessIdentity**
- **cloudfront:CreateDistribution**
- **cloudfront>DeleteCloudFrontOriginAccessIdentity**
- **cloudfront>DeleteDistribution**
- **cloudfront:GetCloudFrontOriginAccessIdentity**
- **cloudfront:GetCloudFrontOriginAccessIdentityConfig**
- **cloudfront:GetDistribution**
- **cloudfront:TagResource**
- **cloudfront:UpdateDistribution**



NOTE

These additional permissions support the use of the **--create-private-s3-bucket** option when processing credentials requests with the **ccoctl aws create-all** command.


Procedure

1. Set a variable for the OpenShift Container Platform release image by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

- Obtain the CCO container image from the OpenShift Container Platform release image by running the following command:

```
$ CCO_IMAGE=$(oc adm release info --image-for='cloud-credential-operator'
$RELEASE_IMAGE -a ~/.pull-secret)
```



NOTE

Ensure that the architecture of the **\$RELEASE_IMAGE** matches the architecture of the environment in which you will use the **ccoctl** tool.

- Extract the **ccoctl** binary from the CCO container image within the OpenShift Container Platform release image by running the following command:

```
$ oc image extract $CCO_IMAGE \
--file="/usr/bin/ccoctl.<rhel_version>" \
-a ~/.pull-secret
```

- For **<rhel_version>**, specify the value that corresponds to the version of Red Hat Enterprise Linux (RHEL) that the host uses. If no value is specified, **ccoctl.rhel8** is used by default. The following values are valid:

- rhel8**: Specify this value for hosts that use RHEL 8.
- rhel9**: Specify this value for hosts that use RHEL 9.

- Change the permissions to make **ccoctl** executable by running the following command:

```
$ chmod 775 ccoctl.<rhel_version>
```

Verification

- To verify that **ccoctl** is ready to use, display the help file by running the following command:

```
$ ccoctl --help
```

Output of **ccoctl --help**

```
OpenShift credentials provisioning tool
```

```
Usage:
```

```
ccoctl [command]
```

```
Available Commands:
```

```
aws      Manage credentials objects for AWS cloud
azure    Manage credentials objects for Azure
gcp      Manage credentials objects for Google cloud
help     Help about any command
ibmcloud Manage credentials objects for IBM Cloud
nutanix  Manage credentials objects for Nutanix
```

Flags:

-h, --help help for ccoctl

Use "ccoctl [command] --help" for more information about a command.

6.3.7.5.2.2. Creating AWS resources with the Cloud Credential Operator utility

You have the following options when creating AWS resources:

- You can use the **ccoctl aws create-all** command to create the AWS resources automatically. This is the quickest way to create the resources. See [Creating AWS resources with a single command](#).
- If you need to review the JSON files that the **ccoctl** tool creates before modifying AWS resources, or if the process the **ccoctl** tool uses to create AWS resources automatically does not meet the requirements of your organization, you can create the AWS resources individually. See [Creating AWS resources individually](#).

6.3.7.5.2.2.1. Creating AWS resources with a single command

If the process the **ccoctl** tool uses to create AWS resources automatically meets the requirements of your organization, you can use the **ccoctl aws create-all** command to automate the creation of AWS resources.

Otherwise, you can create the AWS resources individually. For more information, see "Creating AWS resources individually".



NOTE

By default, **ccoctl** creates objects in the directory in which the commands are run. To create the objects in a different directory, use the **--output-dir** flag. This procedure uses **<path_to_ccoctl_output_dir>** to refer to this directory.

Prerequisites

You must have:

- Extracted and prepared the **ccoctl** binary.

Procedure

1. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

2. Extract the list of **CredentialsRequest** objects from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
```



```
--included \ 1
--install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \ 2
--to=<path_to_directory_for_credentials_requests> 3
```

- 1 The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2 Specify the location of the **install-config.yaml** file.
- 3 Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.




NOTE

This command might take a few moments to run.

3. Use the **ccoctl** tool to process all **CredentialsRequest** objects by running the following command:

```
$ ccoctl aws create-all \
  --name=<name> \ 1
  --region=<aws_region> \ 2
  --credentials-requests-dir=<path_to_credentials_requests_directory> \ 3
  --output-dir=<path_to_ccoctl_output_dir> \ 4
  --create-private-s3-bucket 5
```

- 1 Specify the name used to tag any cloud resources that are created for tracking.
- 2 Specify the AWS region in which cloud resources will be created.
- 3 Specify the directory containing the files for the component **CredentialsRequest** objects.
- 4 Optional: Specify the directory in which you want the **ccoctl** utility to create objects. By default, the utility creates objects in the directory in which the commands are run.
- 5 Optional: By default, the **ccoctl** utility stores the OpenID Connect (OIDC) configuration files in a public S3 bucket and uses the S3 URL as the public OIDC endpoint. To store the OIDC configuration in a private S3 bucket that is accessed by the IAM identity provider through a public CloudFront distribution URL instead, use the **--create-private-s3-bucket** parameter.



NOTE

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

Verification

- To verify that the OpenShift Container Platform secrets are created, list the files in the **<path_to_ccoctl_output_dir>/manifests** directory:

```
$ ls <path_to_ccoctl_output_dir>/manifests
```

Example output

```
cluster-authentication-02-config.yaml
openshift-cloud-credential-operator-cloud-credential-operator-iam-ro-creds-credentials.yaml
openshift-cloud-network-config-controller-cloud-credentials-credentials.yaml
openshift-cluster-api-capamanager-bootstrap-credentials-credentials.yaml
openshift-cluster-csi-drivers-ebs-cloud-credentials-credentials.yaml
openshift-image-registry-installer-cloud-credentials-credentials.yaml
openshift-ingress-operator-cloud-credentials-credentials.yaml
openshift-machine-api-aws-cloud-credentials-credentials.yaml
```

You can verify that the IAM roles are created by querying AWS. For more information, refer to AWS documentation on listing IAM roles.

6.3.7.5.2.2.2. Creating AWS resources individually

You can use the **ccoctl** tool to create AWS resources individually. This option might be useful for an organization that shares the responsibility for creating these resources among different users or departments.

Otherwise, you can use the **ccoctl aws create-all** command to create the AWS resources automatically. For more information, see "Creating AWS resources with a single command".



NOTE

By default, **ccoctl** creates objects in the directory in which the commands are run. To create the objects in a different directory, use the **--output-dir** flag. This procedure uses **<path_to_ccoctl_output_dir>** to refer to this directory.

Some **ccoctl** commands make AWS API calls to create or modify AWS resources. You can use the **--dry-run** flag to avoid making API calls. Using this flag creates JSON files on the local file system instead. You can review and modify the JSON files and then apply them with the AWS CLI tool using the **--cli-input-json** parameters.

Prerequisites

- Extract and prepare the **ccoctl** binary.

Procedure

1. Generate the public and private RSA key files that are used to set up the OpenID Connect provider for the cluster by running the following command:

```
$ ccoctl aws create-key-pair
```

Example output

```
2021/04/13 11:01:02 Generating RSA keypair
2021/04/13 11:01:03 Writing private key to /<path_to_ccoctl_output_dir>/serviceaccount-signer.private
```

```
2021/04/13 11:01:03 Writing public key to /<path_to_ccoctl_output_dir>/serviceaccount-
signer.public
2021/04/13 11:01:03 Copying signing key for use by installer
```

where **serviceaccount-signer.private** and **serviceaccount-signer.public** are the generated key files.

This command also creates a private key that the cluster requires during installation in **/<path_to_ccoctl_output_dir>/tls/bound-service-account-signing-key.key**.

2. Create an OpenID Connect identity provider and S3 bucket on AWS by running the following command:

```
$ ccoctl aws create-identity-provider \
  --name=<name> \ 1
  --region=<aws_region> \ 2
  --public-key-file=<path_to_ccoctl_output_dir>/serviceaccount-signer.public 3
```

- 1** **<name>** is the name used to tag any cloud resources that are created for tracking.
- 2** **<aws-region>** is the AWS region in which cloud resources will be created.
- 3** **<path_to_ccoctl_output_dir>** is the path to the public key file that the **ccoctl aws create-key-pair** command generated.

Example output

```
2021/04/13 11:16:09 Bucket <name>-oidc created
2021/04/13 11:16:10 OpenID Connect discovery document in the S3 bucket <name>-oidc at
.well-known/openid-configuration updated
2021/04/13 11:16:10 Reading public key
2021/04/13 11:16:10 JSON web key set (JWKS) in the S3 bucket <name>-oidc at keys.json
updated
2021/04/13 11:16:18 Identity Provider created with ARN: arn:aws:iam::
<aws_account_id>:oidc-provider/<name>-oidc.s3.<aws_region>.amazonaws.com
```

where **openid-configuration** is a discovery document and **keys.json** is a JSON web key set file.

This command also creates a YAML configuration file in **/<path_to_ccoctl_output_dir>/manifests/cluster-authentication-02-config.yaml**. This file sets the issuer URL field for the service account tokens that the cluster generates, so that the AWS IAM identity provider trusts the tokens.

3. Create IAM roles for each component in the cluster:
 - a. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

- b. Extract the list of **CredentialsRequest** objects from the OpenShift Container Platform release image:

```
$ oc adm release extract \
```

```

--from=$RELEASE_IMAGE \
--credentials-requests \
--included \ 1
--install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \
2
--to=<path_to_directory_for_credentials_requests> 3

```

- 1 The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2 Specify the location of the **install-config.yaml** file.
- 3 Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

- c. Use the **ccoctl** tool to process all **CredentialsRequest** objects by running the following command:

```

$ ccoctl aws create-iam-roles \
  --name=<name> \
  --region=<aws_region> \
  --credentials-requests-dir=<path_to_credentials_requests_directory> \
  --identity-provider-arn=arn:aws:iam::<aws_account_id>:oidc-provider/<name>-oidc.s3.
  <aws_region>.amazonaws.com

```



NOTE

For AWS environments that use alternative IAM API endpoints, such as GovCloud, you must also specify your region with the **--region** parameter.

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

For each **CredentialsRequest** object, **ccoctl** creates an IAM role with a trust policy that is tied to the specified OIDC identity provider, and a permissions policy as defined in each **CredentialsRequest** object from the OpenShift Container Platform release image.

Verification

- To verify that the OpenShift Container Platform secrets are created, list the files in the **<path_to_ccoctl_output_dir>/manifests** directory:

```
$ ls <path_to_ccoctl_output_dir>/manifests
```

Example output

```

cluster-authentication-02-config.yaml
openshift-cloud-credential-operator-cloud-credential-operator-iam-ro-creds-credentials.yaml
openshift-cloud-network-config-controller-cloud-credentials-credentials.yaml
openshift-cluster-api-capamanager-bootstrap-credentials-credentials.yaml
openshift-cluster-csi-drivers-efs-cloud-credentials-credentials.yaml

```

```

openshift-image-registry-installer-cloud-credentials-credentials.yaml
openshift-ingress-operator-cloud-credentials-credentials.yaml
openshift-machine-api-aws-cloud-credentials-credentials.yaml

```

You can verify that the IAM roles are created by querying AWS. For more information, refer to AWS documentation on listing IAM roles.

6.3.7.5.2.3. Incorporating the Cloud Credential Operator utility manifests

To implement short-term security credentials managed outside the cluster for individual components, you must move the manifest files that the Cloud Credential Operator utility (**ccoctl**) created to the correct directories for the installation program.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have configured the Cloud Credential Operator utility (**ccoctl**).
- You have created the cloud provider resources that are required for your cluster with the **ccoctl** utility.

Procedure

1. If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```

apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...

```

2. If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

3. Copy the manifests that the **ccoctl** utility generated to the **manifests** directory that the installation program created by running the following command:

```
$ cp /<path_to_ccoctl_output_dir>/manifests/* ./manifests/
```

4. Copy the private key that the **ccoctl** utility generated in the **tls** directory to the installation directory by running the following command:

```
$ cp -a /<path_to_ccoctl_output_dir>/tls .
```

6.3.7.6. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.

Procedure

1. Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1  
--log-level=info 2
```

- 1** For **<installation_directory>**, specify the location of your customized **./install-config.yaml** file.
- 2** To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

2. Optional: Remove or disable the **AdministratorAccess** policy from the IAM account that you used to install the cluster.



NOTE

The elevated permissions provided by the **AdministratorAccess** policy are required only during installation.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to **<installation_directory>/openshift_install.log**.



IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

6.3.7.7. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

6.3.7.8. Logging in to the cluster by using the web console

The **kubeadmin** user exists by default after an OpenShift Container Platform installation. You can log in to your cluster as the **kubeadmin** user by using the OpenShift Container Platform web console.

Prerequisites

- You have access to the installation host.
- You completed a cluster installation and all cluster Operators are available.

Procedure

1. Obtain the password for the **kubeadmin** user from the **kubeadmin-password** file on the installation host:

```
$ cat <installation_directory>/auth/kubeadmin-password
```



NOTE

Alternatively, you can obtain the **kubeadmin** password from the **<installation_directory>/openshift_install.log** log file on the installation host.

2. List the OpenShift Container Platform web console route:

```
$ oc get routes -n openshift-console | grep 'console-openshift'
```



NOTE

Alternatively, you can obtain the OpenShift Container Platform route from the **<installation_directory>/openshift_install.log** log file on the installation host.

Example output

```
console    console-openshift-console.apps.<cluster_name>.<base_domain>    console
https reencrypt/Redirect None
```

3. Navigate to the route detailed in the output of the preceding command in a web browser and log in as the **kubeadmin** user.

Additional resources

- See [Accessing the web console](#) for more details about accessing and understanding the OpenShift Container Platform web console.

6.3.7.9. Next steps

- [Validating an installation](#).
- [Customize your cluster](#).
- If necessary, you can [opt out of remote health reporting](#) .
- If necessary, you can [remove cloud provider credentials](#).

6.3.8. Installing a cluster on AWS into a government region

In OpenShift Container Platform version 4.16, you can install a cluster on Amazon Web Services (AWS) into a government region. To configure the region, modify parameters in the **install-config.yaml** file before you install the cluster.

6.3.8.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [configured an AWS account](#) to host the cluster.



IMPORTANT

If you have an AWS profile stored on your computer, it must not use a temporary session token that you generated while using a multi-factor authentication device. The cluster continues to use your current AWS credentials to create AWS resources for the entire life of the cluster, so you must use long-term credentials. To generate appropriate keys, see [Managing Access Keys for IAM Users](#) in the AWS documentation. You can supply the keys when you run the installation program.

- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.

6.3.8.2. AWS government regions

OpenShift Container Platform supports deploying a cluster to an [AWS GovCloud \(US\)](#) region.

The following AWS GovCloud partitions are supported:

- **us-gov-east-1**
- **us-gov-west-1**

6.3.8.3. Installation requirements

Before you can install the cluster, you must:

- Provide an existing private AWS VPC and subnets to host the cluster. Public zones are not supported in Route 53 in AWS GovCloud. As a result, clusters must be private when you deploy to an AWS government region.
- Manually create the installation configuration file (**install-config.yaml**).

6.3.8.4. Private clusters

You can deploy a private OpenShift Container Platform cluster that does not expose external endpoints. Private clusters are accessible from only an internal network and are not visible to the internet.



NOTE

Public zones are not supported in Route 53 in an AWS GovCloud Region. Therefore, clusters must be private if they are deployed to an AWS GovCloud Region.

By default, OpenShift Container Platform is provisioned to use publicly-accessible DNS and endpoints. A private cluster sets the DNS, Ingress Controller, and API server to private when you deploy your cluster. This means that the cluster resources are only accessible from your internal network and are not visible to the internet.



IMPORTANT

If the cluster has any public subnets, load balancer services created by administrators might be publicly accessible. To ensure cluster security, verify that these services are explicitly annotated as private.

To deploy a private cluster, you must:

- Use existing networking that meets your requirements. Your cluster resources might be shared between other clusters on the network.
- Deploy from a machine that has access to:
 - The API services for the cloud to which you provision.
 - The hosts on the network that you provision.
 - The internet to obtain installation media.

You can use any machine that meets these access requirements and follows your company's guidelines. For example, this machine can be a bastion host on your cloud network or a machine that has access to the network through a VPN.

6.3.8.4.1. Private clusters in AWS

To create a private cluster on Amazon Web Services (AWS), you must provide an existing private VPC and subnets to host the cluster. The installation program must also be able to resolve the DNS records that the cluster requires. The installation program configures the Ingress Operator and API server for access from only the private network.

The cluster still requires access to internet to access the AWS APIs.

The following items are not required or created when you install a private cluster:

- Public subnets
- Public load balancers, which support public ingress
- A public Route 53 zone that matches the **baseDomain** for the cluster

The installation program does use the **baseDomain** that you specify to create a private Route 53 zone and the required records for the cluster. The cluster is configured so that the Operators do not create public records for the cluster and all cluster machines are placed in the private subnets that you specify.

6.3.8.4.1.1. Limitations

The ability to add public functionality to a private cluster is limited.

- You cannot make the Kubernetes API endpoints public after installation without taking additional actions, including creating public subnets in the VPC for each availability zone in use, creating a public load balancer, and configuring the control plane security groups to allow traffic from the internet on 6443 (Kubernetes API port).
- If you use a public Service type load balancer, you must tag a public subnet in each availability zone with **kubernetes.io/cluster/<cluster-infra-id>: shared** so that AWS can use them to create public load balancers.

6.3.8.5. About using a custom VPC

In OpenShift Container Platform 4.16, you can deploy a cluster into existing subnets in an existing Amazon Virtual Private Cloud (VPC) in Amazon Web Services (AWS). By deploying OpenShift Container Platform into an existing AWS VPC, you might be able to avoid limit constraints in new accounts or more easily abide by the operational constraints that your company's guidelines set. If you cannot obtain the infrastructure creation permissions that are required to create the VPC yourself, use this installation option.

Because the installation program cannot know what other components are also in your existing subnets, it cannot choose subnet CIDRs and so forth on your behalf. You must configure networking for the subnets that you install your cluster to yourself.

6.3.8.5.1. Requirements for using your VPC

The installation program no longer creates the following components:

- Internet gateways
- NAT gateways
- Subnets
- Route tables
- VPCs
- VPC DHCP options
- VPC endpoints



NOTE

The installation program requires that you use the cloud-provided DNS server. Using a custom DNS server is not supported and causes the installation to fail.

If you use a custom VPC, you must correctly configure it and its subnets for the installation program and the cluster to use. See [Amazon VPC console wizard configurations](#) and [Work with VPCs and subnets](#) in the AWS documentation for more information on creating and managing an AWS VPC.

The installation program cannot:

- Subdivide network ranges for the cluster to use.
- Set route tables for the subnets.
- Set VPC options like DHCP.

You must complete these tasks before you install the cluster. See [VPC networking components](#) and [Route tables for your VPC](#) for more information on configuring networking in an AWS VPC.

Your VPC must meet the following characteristics:

- The VPC must not use the **kubernetes.io/cluster/.*: owned, Name**, and **openshift.io/cluster** tags.
The installation program modifies your subnets to add the **kubernetes.io/cluster/.*: shared** tag, so your subnets must have at least one free tag slot available for it. See [Tag Restrictions](#) in the AWS documentation to confirm that the installation program can add a tag to each subnet that you specify. You cannot use a **Name** tag, because it overlaps with the EC2 **Name** field and the installation fails.
- If you want to extend your OpenShift Container Platform cluster into an AWS Outpost and have an existing Outpost subnet, the existing subnet must use the **kubernetes.io/cluster/unmanaged: true** tag. If you do not apply this tag, the installation might fail due to the Cloud Controller Manager creating a service load balancer in the Outpost subnet, which is an unsupported configuration.
- You must enable the **enableDnsSupport** and **enableDnsHostnames** attributes in your VPC, so that the cluster can use the Route 53 zones that are attached to the VPC to resolve cluster's internal DNS records. See [DNS Support in Your VPC](#) in the AWS documentation.
If you prefer to use your own Route 53 hosted private zone, you must associate the existing hosted zone with your VPC prior to installing a cluster. You can define your hosted zone using the **platform.aws.hostedZone** and **platform.aws.hostedZoneRole** fields in the **install-config.yaml** file. You can use a private hosted zone from another account by sharing it with the account where you install the cluster. If you use a private hosted zone from another account, you must use the **Passthrough** or **Manual** credentials mode.

If you are working in a disconnected environment, you are unable to reach the public IP addresses for EC2, ELB, and S3 endpoints. Depending on the level to which you want to restrict internet traffic during the installation, the following configuration options are available:

Option 1: Create VPC endpoints

Create a VPC endpoint and attach it to the subnets that the clusters are using. Name the endpoints as follows:

- **ec2.<aws_region>.amazonaws.com**
- **elasticloadbalancing.<aws_region>.amazonaws.com**
- **s3.<aws_region>.amazonaws.com**

With this option, network traffic remains private between your VPC and the required AWS services.

Option 2: Create a proxy without VPC endpoints

As part of the installation process, you can configure an HTTP or HTTPS proxy. With this option, internet traffic goes through the proxy to reach the required AWS services.

Option 3: Create a proxy with VPC endpoints

As part of the installation process, you can configure an HTTP or HTTPS proxy with VPC endpoints. Create a VPC endpoint and attach it to the subnets that the clusters are using. Name the endpoints as follows:

- **ec2.<aws_region>.amazonaws.com**
- **elasticloadbalancing.<aws_region>.amazonaws.com**
- **s3.<aws_region>.amazonaws.com**

When configuring the proxy in the **install-config.yaml** file, add these endpoints to the **noProxy** field. With this option, the proxy prevents the cluster from accessing the internet directly. However, network traffic remains private between your VPC and the required AWS services.

Required VPC components

You must provide a suitable VPC and subnets that allow communication to your machines.

Component	AWS type	Description
VPC	<ul style="list-style-type: none"> • AWS::EC2::VPC • AWS::EC2::VPCEndpoint 	You must provide a public VPC for the cluster to use. The VPC uses an endpoint that references the route tables for each subnet to improve communication with the registry that is hosted in S3.
Public subnets	<ul style="list-style-type: none"> • AWS::EC2::Subnet • AWS::EC2::SubnetNetworkACLAssociation 	Your VPC must have public subnets for between 1 and 3 availability zones and associate them with appropriate Ingress rules.
Internet gateway	<ul style="list-style-type: none"> • AWS::EC2::InternetGateway • AWS::EC2::VPCGatewayAttachment • AWS::EC2::RouteTable • AWS::EC2::Route • AWS::EC2::SubnetRouteTableAssociation • AWS::EC2::NatGateway • AWS::EC2::EIP 	You must have a public internet gateway, with public routes, attached to the VPC. In the provided templates, each public subnet has a NAT gateway with an EIP address. These NAT gateways allow cluster resources, like private subnet instances, to reach the internet and are not required for some restricted network or proxy scenarios.

Component	AWS type	Description	
Network access control	<ul style="list-style-type: none"> ● AWS::EC2::NetworkAcl ● AWS::EC2::NetworkAclEntry 	You must allow the VPC to access the following ports:	
		Port	Reason
		80	Inbound HTTP traffic
		443	Inbound HTTPS traffic
		22	Inbound SSH traffic
		1024 - 65535	Inbound ephemeral traffic
Private subnets	<ul style="list-style-type: none"> ● AWS::EC2::Subnet ● AWS::EC2::RouteTable ● AWS::EC2::SubnetRouteTableAssociation 	Your VPC can have private subnets. The provided CloudFormation templates can create private subnets for between 1 and 3 availability zones. If you use private subnets, you must provide appropriate routes and tables for them.	
		Port	Reason
		80	Inbound HTTP traffic
		443	Inbound HTTPS traffic
		22	Inbound SSH traffic
		1024 - 65535	Inbound ephemeral traffic
0 - 65535	Outbound ephemeral traffic		

6.3.8.5.2. VPC validation

To ensure that the subnets that you provide are suitable, the installation program confirms the following data:

- All the subnets that you specify exist.
- You provide private subnets.
- The subnet CIDRs belong to the machine CIDR that you specified.
- You provide subnets for each availability zone. Each availability zone contains no more than one public and one private subnet. If you use a private cluster, provide only a private subnet for each availability zone. Otherwise, provide exactly one public and private subnet for each availability zone.
- You provide a public subnet for each private subnet availability zone. Machines are not provisioned in availability zones that you do not provide private subnets for.

If you destroy a cluster that uses an existing VPC, the VPC is not deleted. When you remove the OpenShift Container Platform cluster from a VPC, the **kubernetes.io/cluster/.*: shared** tag is removed from the subnets that it used.

6.3.8.5.3. Division of permissions

Starting with OpenShift Container Platform 4.3, you do not need all of the permissions that are required for an installation program-provisioned infrastructure cluster to deploy a cluster. This change mimics the division of permissions that you might have at your company: some individuals can create different resource in your clouds than others. For example, you might be able to create application-specific items, like instances, buckets, and load balancers, but not networking-related components such as VPCs, subnets, or ingress rules.

The AWS credentials that you use when you create your cluster do not need the networking permissions that are required to make VPCs and core networking components within the VPC, such as subnets, routing tables, internet gateways, NAT, and VPN. You still need permission to make the application resources that the machines within the cluster require, such as ELBs, security groups, S3 buckets, and nodes.

6.3.8.5.4. Isolation between clusters

If you deploy OpenShift Container Platform to an existing network, the isolation of cluster services is reduced in the following ways:

- You can install multiple OpenShift Container Platform clusters in the same VPC.
- ICMP ingress is allowed from the entire network.
- TCP 22 ingress (SSH) is allowed to the entire network.
- Control plane TCP 6443 ingress (Kubernetes API) is allowed to the entire network.
- Control plane TCP 22623 ingress (MCS) is allowed to the entire network.

6.3.8.5.5. Optional: AWS security groups

By default, the installation program creates and attaches security groups to control plane and compute machines. The rules associated with the default security groups cannot be modified.

However, you can apply additional existing AWS security groups, which are associated with your existing VPC, to control plane and compute machines. Applying custom security groups can help you meet the security needs of your organization, in such cases where you need to control the incoming or outgoing traffic of these machines.

As part of the installation process, you apply custom security groups by modifying the **install-config.yaml** file before deploying the cluster.

For more information, see "Applying existing AWS security groups to the cluster".

6.3.8.6. Obtaining an AWS Marketplace image

If you are deploying an OpenShift Container Platform cluster using an AWS Marketplace image, you must first subscribe through AWS. Subscribing to the offer provides you with the AMI ID that the installation program uses to deploy compute nodes.

Prerequisites

- You have an AWS account to purchase the offer. This account does not have to be the same account that is used to install the cluster.

Procedure

1. Complete the OpenShift Container Platform subscription from the [AWS Marketplace](#).
2. Record the AMI ID for your specific AWS Region. As part of the installation process, you must update the **install-config.yaml** file with this value before deploying the cluster.

Sample **install-config.yaml** file with AWS Marketplace compute nodes

```
apiVersion: v1
baseDomain: example.com
compute:
- hyperthreading: Enabled
  name: worker
  platform:
    aws:
      amiID: ami-06c4d345f7c207239 1
      type: m5.4xlarge
      replicas: 3
  metadata:
    name: test-cluster
  platform:
    aws:
      region: us-east-2 2
  sshKey: ssh-ed25519 AAAA...
  pullSecret: '{"auths": ...}'
```

- 1** The AMI ID from your AWS Marketplace subscription.
- 2** Your AMI ID is associated with a specific AWS Region. When creating the installation configuration file, ensure that you select the same AWS Region that you specified when configuring your subscription.

6.3.8.7. Manually creating the installation configuration file

Installing the cluster requires that you manually create the installation configuration file.


Prerequisites

- You have an SSH public key on your local machine to provide to the installation program. The key will be used for SSH authentication onto your cluster nodes for debugging and disaster recovery.
- You have obtained the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create an installation directory to store your required installation assets in:


```
$ mkdir <installation_directory>
```



IMPORTANT

You must create a directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.


2. Customize the sample **install-config.yaml** file template that is provided and save it in the **<installation_directory>**.



NOTE

You must name this configuration file **install-config.yaml**.

3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the next step of the installation process. You must back it up now.

Additional resources

- [Installation configuration parameters for AWS](#)

6.3.8.7.1. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 6.17. Minimum resource requirements

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Bootstrap	RHCOS	4	16 GB	100 GB	300
Control plane	RHCOS	4	16 GB	100 GB	300
Compute	RHCOS, RHEL 8.6 and later [3]	2	8 GB	100 GB	300

1. One vCPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio: (threads per core × cores) × sockets = vCPUs.

2. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes which require a 10 ms p99 fsync duration. Note that on many cloud platforms, storage size and IOPS scale together, so you might need to over-allocate storage volume to obtain sufficient performance.
3. As with all user-provisioned installations, if you choose to use RHEL compute machines in your cluster, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Use of RHEL 7 compute machines is deprecated and has been removed in OpenShift Container Platform 4.10 and later.



NOTE

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

6.3.8.7.2. Tested instance types for AWS

The following Amazon Web Services (AWS) instance types have been tested with OpenShift Container Platform.



NOTE

Use the machine types included in the following charts for your AWS instances. If you use an instance type that is not listed in the chart, ensure that the instance size you use matches the minimum resource requirements that are listed in the section named "Minimum resource requirements for cluster installation".

Example 6.37. Machine types based on 64-bit x86 architecture

- **c4.***
- **c5.***
- **c5a.***
- **i3.***

- m4.*
- m5.*
- m5a.*
- m6a.*
- m6i.*
- r4.*
- r5.*
- r5a.*
- r6i.*
- t3.*
- t3a.*

6.3.8.7.3. Tested instance types for AWS on 64-bit ARM infrastructures

The following Amazon Web Services (AWS) 64-bit ARM instance types have been tested with OpenShift Container Platform.



NOTE

Use the machine types included in the following charts for your AWS ARM instances. If you use an instance type that is not listed in the chart, ensure that the instance size you use matches the minimum resource requirements that are listed in "Minimum resource requirements for cluster installation".

Example 6.38. Machine types based on 64-bit ARM architecture

- c6g.*
- m6g.*

6.3.8.7.4. Sample customized install-config.yaml file for AWS

You can customize the installation configuration file (**install-config.yaml**) to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.



IMPORTANT

This sample YAML file is provided for reference only. Use it as a resource to enter parameter values into the installation configuration file that you created manually.

```
apiVersion: v1
baseDomain: example.com 1
```

```
credentialsMode: Mint 2
controlPlane: 3 4
  hyperthreading: Enabled 5
  name: master
  platform:
    aws:
      zones:
        - us-gov-west-1a
        - us-gov-west-1b
      rootVolume:
        iops: 4000
        size: 500
        type: io1 6
      metadataService:
        authentication: Optional 7
        type: m6i.xlarge
      replicas: 3
compute: 8
  - hyperthreading: Enabled 9
  name: worker
  platform:
    aws:
      rootVolume:
        iops: 2000
        size: 500
        type: io1 10
      metadataService:
        authentication: Optional 11
        type: c5.4xlarge
      zones:
        - us-gov-west-1c
      replicas: 3
metadata:
  name: test-cluster 12
networking:
  clusterNetwork:
    - cidr: 10.128.0.0/14
    hostPrefix: 23
  machineNetwork:
    - cidr: 10.0.0.0/16
  networkType: OVNKubernetes 13
  serviceNetwork:
    - 172.30.0.0/16
platform:
  aws:
    region: us-gov-west-1 14
    propagateUserTags: true 15
    userTags:
      adminContact: jdoe
      costCenter: 7536
  subnets: 16
    - subnet-1
    - subnet-2
    - subnet-3
```

```

amiID: ami-0c5d3e03c0ab9b19a 17
serviceEndpoints: 18
  - name: ec2
    url: https://vpce-id.ec2.us-west-2.vpce.amazonaws.com
hostedZone: Z3URY6TWQ91KVV 19
fips: false 20
sshKey: ssh-ed25519 AAAA... 21
publish: Internal 22
pullSecret: '{"auths": ...}' 23

```


1 12 14 23 Required.

2 Optional: Add this parameter to force the Cloud Credential Operator (CCO) to use the specified mode. By default, the CCO uses the root credentials in the **kube-system** namespace to dynamically try to determine the capabilities of the credentials. For details about CCO modes, see the "About the Cloud Credential Operator" section in the *Authentication and authorization* guide.

3 8 15 If you do not provide these parameters and values, the installation program provides the default value.

4 The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.

5 9 Whether to enable or disable simultaneous multithreading, or **hyperthreading**. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores. You can disable it by setting the parameter value to **Disabled**. If you disable simultaneous multithreading in some cluster machines, you must disable it in all cluster machines.



IMPORTANT

If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance. Use larger instance types, such as **m4.2xlarge** or **m5.2xlarge**, for your machines if you disable simultaneous multithreading.

6 10 To configure faster storage for etcd, especially for larger clusters, set the storage type as **io1** and set **iops** to **2000**.

7 11 Whether to require the [Amazon EC2 Instance Metadata Service v2](#) (IMDSv2). To require IMDSv2, set the parameter value to **Required**. To allow the use of both IMDSv1 and IMDSv2, set the parameter value to **Optional**. If no value is specified, both IMDSv1 and IMDSv2 are allowed.



NOTE

The IMDS configuration for control plane machines that is set during cluster installation can only be changed by using the AWS CLI. The IMDS configuration for compute machines can be changed by using compute machine sets.

13 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.

16 If you provide your own VPC, specify subnets for each availability zone that your cluster uses.

- 17 The ID of the AMI used to boot machines for the cluster. If set, the AMI must belong to the same region as the cluster.
- 18 The AWS service endpoints. Custom endpoints are required when installing to an unknown AWS region. The endpoint URL must use the **https** protocol and the host must trust the certificate.
- 19 The ID of your existing Route 53 private hosted zone. Providing an existing hosted zone requires that you supply your own VPC and the hosted zone is already associated with the VPC prior to installing your cluster. If undefined, the installation program creates a new hosted zone.
- 20 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.



IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 21 You can optionally provide the **sshKey** value that you use to access the machines in your cluster.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- 22 How to publish the user-facing endpoints of your cluster. Set **publish** to **Internal** to deploy a private cluster, which cannot be accessed from the internet. The default value is **External**.

6.3.8.7.5. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: ec2.<aws_region>.amazonaws.com,elasticloadbalancing.
  <aws_region>.amazonaws.com,s3.<aws_region>.amazonaws.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations. If you have added the Amazon **EC2,Elastic Load Balancing**, and **S3** VPC endpoints to your VPC, you must add these endpoints to the **noProxy** field.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.

**NOTE**

The installation program does not support the proxy **readinessEndpoints** field.

**NOTE**

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.

**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

6.3.8.7.6. Applying existing AWS security groups to the cluster

Applying existing AWS security groups to your control plane and compute machines can help you meet the security needs of your organization, in such cases where you need to control the incoming or outgoing traffic of these machines.

Prerequisites

- You have created the security groups in AWS. For more information, see the AWS documentation about working with [security groups](#).
- The security groups must be associated with the existing VPC that you are deploying the cluster to. The security groups cannot be associated with another VPC.
- You have an existing **install-config.yaml** file.

Procedure

1. In the **install-config.yaml** file, edit the **compute.platform.aws.additionalSecurityGroupIDs** parameter to specify one or more custom security groups for your compute machines.
2. Edit the **controlPlane.platform.aws.additionalSecurityGroupIDs** parameter to specify one or more custom security groups for your control plane machines.
3. Save the file and reference it when deploying the cluster.

Sample **install-config.yaml** file that specifies custom security groups

```
# ...
compute:
- hyperthreading: Enabled
  name: worker
```



```

platform:
  aws:
    additionalSecurityGroupIDs:
      - sg-1 ❶
      - sg-2
  replicas: 3
controlPlane:
  hyperthreading: Enabled
  name: master
platform:
  aws:
    additionalSecurityGroupIDs:
      - sg-3
      - sg-4
  replicas: 3
platform:
  aws:
    region: us-east-1
  subnets: ❷
    - subnet-1
    - subnet-2
    - subnet-3

```

- ❶ Specify the name of the security group as it appears in the Amazon EC2 console, including the **sg** prefix.
- ❷ Specify subnets for each availability zone that your cluster uses.

6.3.8.8. Alternatives to storing administrator-level secrets in the kube-system project

By default, administrator secrets are stored in the **kube-system** project. If you configured the **credentialsMode** parameter in the **install-config.yaml** file to **Manual**, you must use one of the following alternatives:

- To manage long-term cloud credentials manually, follow the procedure in [Manually creating long-term credentials](#).
- To implement short-term credentials that are managed outside the cluster for individual components, follow the procedures in [Incorporating the Cloud Credential Operator utility manifests](#).

6.3.8.8.1. Manually creating long-term credentials

The Cloud Credential Operator (CCO) can be put into manual mode prior to installation in environments where the cloud identity and access management (IAM) APIs are not reachable, or the administrator prefers not to store an administrator-level credential secret in the cluster **kube-system** namespace.

Procedure

1. If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```

apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...

```

- If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

- Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

- Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```

$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included 1 \
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml 2 \
  --to=<path_to_directory_for_credentials_requests> 3

```

- The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- Specify the location of the **install-config.yaml** file.
- Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

This command creates a YAML file for each **CredentialsRequest** object.

Sample CredentialsRequest object

```

apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: AWSProviderSpec
    statementEntries:
      - effect: Allow
        action:
          - iam:GetUser

```

```

- iam:GetUserPolicy
- iam:ListAccessKeys
resource: "*"
...

```

5. Create YAML files for secrets in the **openshift-install** manifests directory that you generated previously. The secrets must be stored using the namespace and secret name defined in the **spec.secretRef** for each **CredentialsRequest** object.

Sample CredentialsRequest object with secrets

```

apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: AWSProviderSpec
    statementEntries:
      - effect: Allow
        action:
          - s3:CreateBucket
          - s3>DeleteBucket
        resource: "*"
    ...
  secretRef:
    name: <component_secret>
    namespace: <component_namespace>
...

```

Sample Secret object

```

apiVersion: v1
kind: Secret
metadata:
  name: <component_secret>
  namespace: <component_namespace>
data:
  aws_access_key_id: <base64_encoded_aws_access_key_id>
  aws_secret_access_key: <base64_encoded_aws_secret_access_key>

```



IMPORTANT

Before upgrading a cluster that uses manually maintained credentials, you must ensure that the CCO is in an upgradeable state.

6.3.8.8.2. Configuring an AWS cluster to use short-term credentials

To install a cluster that is configured to use the AWS Security Token Service (STS), you must configure the CCO utility and create the required AWS resources for your cluster.

6.3.8.8.2.1. Configuring the Cloud Credential Operator utility

To create and manage cloud credentials from outside of the cluster when the Cloud Credential Operator (CCO) is operating in manual mode, extract and prepare the CCO utility (**ccoctl**) binary.



NOTE

The **ccoctl** utility is a Linux binary that must run in a Linux environment.

Prerequisites

- You have access to an OpenShift Container Platform account with cluster administrator access.
- You have installed the OpenShift CLI (**oc**).
- You have created an AWS account for the **ccoctl** utility to use with the following permissions:

Example 6.39. Required AWS permissions

Required iam permissions

- **iam:CreateOpenIDConnectProvider**
- **iam:CreateRole**
- **iam>DeleteOpenIDConnectProvider**
- **iam>DeleteRole**
- **iam>DeleteRolePolicy**
- **iam:GetOpenIDConnectProvider**
- **iam:GetRole**
- **iam:GetUser**
- **iam:ListOpenIDConnectProviders**
- **iam:ListRolePolicies**
- **iam:ListRoles**
- **iam:PutRolePolicy**
- **iam:TagOpenIDConnectProvider**
- **iam:TagRole**

Required s3 permissions

- **s3:CreateBucket**
- **s3>DeleteBucket**
- **s3>DeleteObject**
- **s3:GetBucketAcl**

- **s3:GetBucketTagging**
- **s3:GetObject**
- **s3:GetObjectAcl**
- **s3:GetObjectTagging**
- **s3:ListBucket**
- **s3:PutBucketAcl**
- **s3:PutBucketPolicy**
- **s3:PutBucketPublicAccessBlock**
- **s3:PutBucketTagging**
- **s3:PutObject**
- **s3:PutObjectAcl**
- **s3:PutObjectTagging**


Required cloudfront permissions

- **cloudfront:ListCloudFrontOriginAccessIdentities**
- **cloudfront:ListDistributions**
- **cloudfront:ListTagsForResource**

If you plan to store the OIDC configuration in a private S3 bucket that is accessed by the IAM identity provider through a public CloudFront distribution URL, the AWS account that runs the **ccoctl** utility requires the following additional permissions:

Example 6.40. Additional permissions for a private S3 bucket with CloudFront

- **cloudfront:CreateCloudFrontOriginAccessIdentity**
- **cloudfront:CreateDistribution**
- **cloudfront>DeleteCloudFrontOriginAccessIdentity**
- **cloudfront>DeleteDistribution**
- **cloudfront:GetCloudFrontOriginAccessIdentity**
- **cloudfront:GetCloudFrontOriginAccessIdentityConfig**
- **cloudfront:GetDistribution**
- **cloudfront:TagResource**
- **cloudfront:UpdateDistribution**

**NOTE**

These additional permissions support the use of the **--create-private-s3-bucket** option when processing credentials requests with the **ccoctl aws create-all** command.

Procedure

1. Set a variable for the OpenShift Container Platform release image by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

2. Obtain the CCO container image from the OpenShift Container Platform release image by running the following command:

```
$ CCO_IMAGE=$(oc adm release info --image-for='cloud-credential-operator'
$RELEASE_IMAGE -a ~/.pull-secret)
```

**NOTE**

Ensure that the architecture of the **\$RELEASE_IMAGE** matches the architecture of the environment in which you will use the **ccoctl** tool.

3. Extract the **ccoctl** binary from the CCO container image within the OpenShift Container Platform release image by running the following command:

```
$ oc image extract $CCO_IMAGE \
--file="/usr/bin/ccoctl.<rhel_version>" \
-a ~/.pull-secret
```

- 1 For **<rhel_version>**, specify the value that corresponds to the version of Red Hat Enterprise Linux (RHEL) that the host uses. If no value is specified, **ccoctl.rhel8** is used by default. The following values are valid:

- **rhel8**: Specify this value for hosts that use RHEL 8.
- **rhel9**: Specify this value for hosts that use RHEL 9.

4. Change the permissions to make **ccoctl** executable by running the following command:

```
$ chmod 775 ccoctl.<rhel_version>
```

Verification

- To verify that **ccoctl** is ready to use, display the help file by running the following command:

```
$ ccoctl --help
```

Output of ccoctl --help

OpenShift credentials provisioning tool

Usage:

`ccoctl [command]`

Available Commands:

`aws` Manage credentials objects for AWS cloud
`azure` Manage credentials objects for Azure
`gcp` Manage credentials objects for Google cloud
`help` Help about any command
`ibmcloud` Manage credentials objects for IBM Cloud
`nutanix` Manage credentials objects for Nutanix

Flags:

`-h, --help` help for `ccoctl`

Use "`ccoctl [command] --help`" for more information about a command.

6.3.8.8.2.2. Creating AWS resources with the Cloud Credential Operator utility

You have the following options when creating AWS resources:

- You can use the **`ccoctl aws create-all`** command to create the AWS resources automatically. This is the quickest way to create the resources. See [Creating AWS resources with a single command](#).
- If you need to review the JSON files that the **`ccoctl`** tool creates before modifying AWS resources, or if the process the **`ccoctl`** tool uses to create AWS resources automatically does not meet the requirements of your organization, you can create the AWS resources individually. See [Creating AWS resources individually](#).

6.3.8.8.2.2.1. Creating AWS resources with a single command

If the process the **`ccoctl`** tool uses to create AWS resources automatically meets the requirements of your organization, you can use the **`ccoctl aws create-all`** command to automate the creation of AWS resources.

Otherwise, you can create the AWS resources individually. For more information, see "Creating AWS resources individually".



NOTE

By default, **`ccoctl`** creates objects in the directory in which the commands are run. To create the objects in a different directory, use the **`--output-dir`** flag. This procedure uses **`<path_to_ccoctl_output_dir>`** to refer to this directory.

Prerequisites

You must have:

- Extracted and prepared the **`ccoctl`** binary.

Procedure

1. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

2. Extract the list of **CredentialsRequest** objects from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included \1
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \2
  --to=<path_to_directory_for_credentials_requests> \3
```

- 1 The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2 Specify the location of the **install-config.yaml** file.
- 3 Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.




NOTE

This command might take a few moments to run.

3. Use the **ccoctl** tool to process all **CredentialsRequest** objects by running the following command:

```
$ ccoctl aws create-all \
  --name=<name> \1
  --region=<aws_region> \2
  --credentials-requests-dir=<path_to_credentials_requests_directory> \3
  --output-dir=<path_to_ccoctl_output_dir> \4
  --create-private-s3-bucket \5
```

- 1 Specify the name used to tag any cloud resources that are created for tracking.
- 2 Specify the AWS region in which cloud resources will be created.
- 3 Specify the directory containing the files for the component **CredentialsRequest** objects.
- 4 Optional: Specify the directory in which you want the **ccoctl** utility to create objects. By default, the utility creates objects in the directory in which the commands are run.
- 5 Optional: By default, the **ccoctl** utility stores the OpenID Connect (OIDC) configuration files in a public S3 bucket and uses the S3 URL as the public OIDC endpoint. To store the OIDC configuration in a private S3 bucket that is accessed by the IAM identity provider through a public CloudFront distribution URL instead, use the **--create-private-s3-bucket** parameter.

**NOTE**

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

Verification

- To verify that the OpenShift Container Platform secrets are created, list the files in the **<path_to_ccoctl_output_dir>/manifests** directory:

```
$ ls <path_to_ccoctl_output_dir>/manifests
```

Example output

```
cluster-authentication-02-config.yaml
openshift-cloud-credential-operator-cloud-credential-operator-iam-ro-creds-credentials.yaml
openshift-cloud-network-config-controller-cloud-credentials-credentials.yaml
openshift-cluster-api-capamanager-bootstrap-credentials-credentials.yaml
openshift-cluster-csi-drivers-ebs-cloud-credentials-credentials.yaml
openshift-image-registry-installer-cloud-credentials-credentials.yaml
openshift-ingress-operator-cloud-credentials-credentials.yaml
openshift-machine-api-aws-cloud-credentials-credentials.yaml
```

You can verify that the IAM roles are created by querying AWS. For more information, refer to AWS documentation on listing IAM roles.

6.3.8.8.2.2.2. Creating AWS resources individually

You can use the **ccoctl** tool to create AWS resources individually. This option might be useful for an organization that shares the responsibility for creating these resources among different users or departments.

Otherwise, you can use the **ccoctl aws create-all** command to create the AWS resources automatically. For more information, see "Creating AWS resources with a single command".

**NOTE**

By default, **ccoctl** creates objects in the directory in which the commands are run. To create the objects in a different directory, use the **--output-dir** flag. This procedure uses **<path_to_ccoctl_output_dir>** to refer to this directory.

Some **ccoctl** commands make AWS API calls to create or modify AWS resources. You can use the **--dry-run** flag to avoid making API calls. Using this flag creates JSON files on the local file system instead. You can review and modify the JSON files and then apply them with the AWS CLI tool using the **--cli-input-json** parameters.

Prerequisites

- Extract and prepare the **ccoctl** binary.

Procedure

1. Generate the public and private RSA key files that are used to set up the OpenID Connect provider for the cluster by running the following command:

```
$ ccoctl aws create-key-pair
```

Example output

```
2021/04/13 11:01:02 Generating RSA keypair
2021/04/13 11:01:03 Writing private key to /<path_to_ccoctl_output_dir>/serviceaccount-
signer.private
2021/04/13 11:01:03 Writing public key to /<path_to_ccoctl_output_dir>/serviceaccount-
signer.public
2021/04/13 11:01:03 Copying signing key for use by installer
```

where **serviceaccount-signer.private** and **serviceaccount-signer.public** are the generated key files.

This command also creates a private key that the cluster requires during installation in **/<path_to_ccoctl_output_dir>/tls/bound-service-account-signing-key.key**.

2. Create an OpenID Connect identity provider and S3 bucket on AWS by running the following command:

```
$ ccoctl aws create-identity-provider \
  --name=<name> \ 1
  --region=<aws_region> \ 2
  --public-key-file=<path_to_ccoctl_output_dir>/serviceaccount-signer.public 3
```

- 1** **<name>** is the name used to tag any cloud resources that are created for tracking.
- 2** **<aws-region>** is the AWS region in which cloud resources will be created.
- 3** **<path_to_ccoctl_output_dir>** is the path to the public key file that the **ccoctl aws create-key-pair** command generated.

Example output

```
2021/04/13 11:16:09 Bucket <name>-oidc created
2021/04/13 11:16:10 OpenID Connect discovery document in the S3 bucket <name>-oidc at
.well-known/openid-configuration updated
2021/04/13 11:16:10 Reading public key
2021/04/13 11:16:10 JSON web key set (JWKS) in the S3 bucket <name>-oidc at keys.json
updated
2021/04/13 11:16:18 Identity Provider created with ARN: arn:aws:iam::
<aws_account_id>:oidc-provider/<name>-oidc.s3.<aws_region>.amazonaws.com
```

where **openid-configuration** is a discovery document and **keys.json** is a JSON web key set file.

This command also creates a YAML configuration file in **/<path_to_ccoctl_output_dir>/manifests/cluster-authentication-02-config.yaml**. This file sets the issuer URL field for the service account tokens that the cluster generates, so that the AWS IAM identity provider trusts the tokens.

3. Create IAM roles for each component in the cluster:

- a. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

- b. Extract the list of **CredentialsRequest** objects from the OpenShift Container Platform release image:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included 1 \
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \
2 \
  --to=<path_to_directory_for_credentials_requests> 3
```

- 1** The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2** Specify the location of the **install-config.yaml** file.
- 3** Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

- c. Use the **ccoctl** tool to process all **CredentialsRequest** objects by running the following command:

```
$ ccoctl aws create-iam-roles \
  --name=<name> \
  --region=<aws_region> \
  --credentials-requests-dir=<path_to_credentials_requests_directory> \
  --identity-provider-arn=arn:aws:iam::<aws_account_id>:oidc-provider/<name>-oidc.s3.
  <aws_region>.amazonaws.com
```

**NOTE**

For AWS environments that use alternative IAM API endpoints, such as GovCloud, you must also specify your region with the **--region** parameter.

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

For each **CredentialsRequest** object, **ccoctl** creates an IAM role with a trust policy that is tied to the specified OIDC identity provider, and a permissions policy as defined in each **CredentialsRequest** object from the OpenShift Container Platform release image.

Verification

- To verify that the OpenShift Container Platform secrets are created, list the files in the **<path_to_ccoctl_output_dir>/manifests** directory:

```
$ ls <path_to_ccoctl_output_dir>/manifests
```

Example output

```
cluster-authentication-02-config.yaml
openshift-cloud-credential-operator-cloud-credential-operator-iam-ro-creds-credentials.yaml
openshift-cloud-network-config-controller-cloud-credentials-credentials.yaml
openshift-cluster-api-capamanager-bootstrap-credentials-credentials.yaml
openshift-cluster-csi-drivers-ebs-cloud-credentials-credentials.yaml
openshift-image-registry-installer-cloud-credentials-credentials.yaml
openshift-ingress-operator-cloud-credentials-credentials.yaml
openshift-machine-api-aws-cloud-credentials-credentials.yaml
```

You can verify that the IAM roles are created by querying AWS. For more information, refer to AWS documentation on listing IAM roles.

6.3.8.8.2.3. Incorporating the Cloud Credential Operator utility manifests

To implement short-term security credentials managed outside the cluster for individual components, you must move the manifest files that the Cloud Credential Operator utility (**ccoctl**) created to the correct directories for the installation program.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have configured the Cloud Credential Operator utility (**ccoctl**).
- You have created the cloud provider resources that are required for your cluster with the **ccoctl** utility.

Procedure

1. If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```
apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...
```

2. If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

3. Copy the manifests that the **ccoctl** utility generated to the **manifests** directory that the installation program created by running the following command:

```
$ cp /<path_to_ccoctl_output_dir>/manifests/* ./manifests/
```

4. Copy the private key that the **ccoctl** utility generated in the **tls** directory to the installation directory by running the following command:

```
$ cp -a /<path_to_ccoctl_output_dir>/tls .
```

6.3.8.9. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.

Procedure

1. Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1  
--log-level=info 2
```

- 1** For **<installation_directory>**, specify the location of your customized **./install-config.yaml** file.
- 2** To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

2. Optional: Remove or disable the **AdministratorAccess** policy from the IAM account that you used to install the cluster.



NOTE

The elevated permissions provided by the **AdministratorAccess** policy are required only during installation.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to **<installation_directory>/openshift_install.log**.



IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

6.3.8.10. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

- Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

6.3.8.11. Logging in to the cluster by using the web console

The **kubeadmin** user exists by default after an OpenShift Container Platform installation. You can log in to your cluster as the **kubeadmin** user by using the OpenShift Container Platform web console.

Prerequisites

- You have access to the installation host.
- You completed a cluster installation and all cluster Operators are available.

Procedure

- Obtain the password for the **kubeadmin** user from the **kubeadmin-password** file on the installation host:

```
$ cat <installation_directory>/auth/kubeadmin-password
```



NOTE

Alternatively, you can obtain the **kubeadmin** password from the **<installation_directory>/openshift_install.log** log file on the installation host.

- List the OpenShift Container Platform web console route:

```
$ oc get routes -n openshift-console | grep 'console-openshift'
```



NOTE

Alternatively, you can obtain the OpenShift Container Platform route from the **<installation_directory>/openshift_install.log** log file on the installation host.

Example output

```
console    console-openshift-console.apps.<cluster_name>.<base_domain>    console
https     reencrypt/Redirect    None
```

3. Navigate to the route detailed in the output of the preceding command in a web browser and log in as the **kubeadmin** user.

Additional resources

- See [Accessing the web console](#) for more details about accessing and understanding the OpenShift Container Platform web console.

6.3.8.12. Next steps

- [Validating an installation.](#)
- [Customize your cluster.](#)
- If necessary, you can [opt out of remote health reporting](#) .
- If necessary, you can [remove cloud provider credentials](#).

6.3.9. Installing a cluster on AWS into a Secret or Top Secret Region

In OpenShift Container Platform version 4.16, you can install a cluster on Amazon Web Services (AWS) into the following secret regions:

- Secret Commercial Cloud Services (SC2S)
- Commercial Cloud Services (C2S)

To configure a cluster in either region, you change parameters in the **install config.yaml** file before you install the cluster.



WARNING

In OpenShift Container Platform 4.16, the installation program uses Cluster API instead of Terraform to provision cluster infrastructure during installations on AWS. Installing a cluster on AWS into a secret or top-secret region by using the Cluster API implementation has not been tested as of the release of OpenShift Container Platform 4.16. This document will be updated when installation into a secret region has been tested.

There is a known issue with Network Load Balancers' support for security groups in secret or top secret regions that causes installations in these regions to fail. For more information, see [OCPBUGS-33311](#).

6.3.9.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).

- You [configured an AWS account](#) to host the cluster.



IMPORTANT

If you have an AWS profile stored on your computer, it must not use a temporary session token that you generated while using a multifactor authentication device. The cluster continues to use your current AWS credentials to create AWS resources for the entire life of the cluster, so you must use long-term credentials. To generate appropriate keys, see [Managing Access Keys for IAM Users](#) in the AWS documentation. You can supply the keys when you run the installation program.

- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.

6.3.9.2. AWS secret regions

The following AWS secret partitions are supported:

- **us-isob-east-1** (SC2S)
- **us-iso-east-1** (C2S)



NOTE

The maximum supported MTU in an AWS SC2S and C2S Regions is not the same as AWS commercial. For more information about configuring MTU during installation, see the *Cluster Network Operator configuration object* section in *Installing a cluster on AWS with network customizations*

6.3.9.3. Installation requirements

Red Hat does not publish a Red Hat Enterprise Linux CoreOS (RHCOS) Amazon Machine Image for the AWS Secret and Top Secret Regions.

Before you can install the cluster, you must:

- Upload a custom RHCOS AMI.
- Manually create the installation configuration file (**install-config.yaml**).
- Specify the AWS region, and the accompanying custom AMI, in the installation configuration file.

You cannot use the OpenShift Container Platform installation program to create the installation configuration file. The installer does not list an AWS region without native support for an RHCOS AMI.



IMPORTANT

You must also define a custom CA certificate in the **additionalTrustBundle** field of the **install-config.yaml** file because the AWS API requires a custom CA trust bundle. To allow the installation program to access the AWS API, the CA certificates must also be defined on the machine that runs the installation program. You must add the CA bundle to the trust store on the machine, use the **AWS_CA_BUNDLE** environment variable, or define the CA bundle in the **ca_bundle** field of the AWS config file.

6.3.9.4. Private clusters

You can deploy a private OpenShift Container Platform cluster that does not expose external endpoints. Private clusters are accessible from only an internal network and are not visible to the internet.



NOTE

Public zones are not supported in Route 53 in an AWS Top Secret Region. Therefore, clusters must be private if they are deployed to an AWS Top Secret Region.

By default, OpenShift Container Platform is provisioned to use publicly-accessible DNS and endpoints. A private cluster sets the DNS, Ingress Controller, and API server to private when you deploy your cluster. This means that the cluster resources are only accessible from your internal network and are not visible to the internet.



IMPORTANT

If the cluster has any public subnets, load balancer services created by administrators might be publicly accessible. To ensure cluster security, verify that these services are explicitly annotated as private.

To deploy a private cluster, you must:

- Use existing networking that meets your requirements. Your cluster resources might be shared between other clusters on the network.
- Deploy from a machine that has access to:
 - The API services for the cloud to which you provision.
 - The hosts on the network that you provision.
 - The internet to obtain installation media.

You can use any machine that meets these access requirements and follows your company's guidelines. For example, this machine can be a bastion host on your cloud network or a machine that has access to the network through a VPN.

6.3.9.4.1. Private clusters in AWS

To create a private cluster on Amazon Web Services (AWS), you must provide an existing private VPC and subnets to host the cluster. The installation program must also be able to resolve the DNS records that the cluster requires. The installation program configures the Ingress Operator and API server for access from only the private network.

The cluster still requires access to internet to access the AWS APIs.

The following items are not required or created when you install a private cluster:

- Public subnets
- Public load balancers, which support public ingress
- A public Route 53 zone that matches the **baseDomain** for the cluster

The installation program does use the **baseDomain** that you specify to create a private Route 53 zone and the required records for the cluster. The cluster is configured so that the Operators do not create public records for the cluster and all cluster machines are placed in the private subnets that you specify.

6.3.9.4.1.1. Limitations

The ability to add public functionality to a private cluster is limited.

- You cannot make the Kubernetes API endpoints public after installation without taking additional actions, including creating public subnets in the VPC for each availability zone in use, creating a public load balancer, and configuring the control plane security groups to allow traffic from the internet on 6443 (Kubernetes API port).
- If you use a public Service type load balancer, you must tag a public subnet in each availability zone with **kubernetes.io/cluster/<cluster-infra-id>: shared** so that AWS can use them to create public load balancers.

6.3.9.5. About using a custom VPC

In OpenShift Container Platform 4.16, you can deploy a cluster into existing subnets in an existing Amazon Virtual Private Cloud (VPC) in Amazon Web Services (AWS). By deploying OpenShift Container Platform into an existing AWS VPC, you might be able to avoid limit constraints in new accounts or more easily abide by the operational constraints that your company's guidelines set. If you cannot obtain the infrastructure creation permissions that are required to create the VPC yourself, use this installation option.

Because the installation program cannot know what other components are also in your existing subnets, it cannot choose subnet CIDRs and so forth on your behalf. You must configure networking for the subnets that you install your cluster to yourself.

6.3.9.5.1. Requirements for using your VPC

The installation program no longer creates the following components:

- Internet gateways
- NAT gateways
- Subnets
- Route tables
- VPCs
- VPC DHCP options
- VPC endpoints



NOTE

The installation program requires that you use the cloud-provided DNS server. Using a custom DNS server is not supported and causes the installation to fail.

If you use a custom VPC, you must correctly configure it and its subnets for the installation program and the cluster to use. See [Amazon VPC console wizard configurations](#) and [Work with VPCs and subnets](#) in the AWS documentation for more information on creating and managing an AWS VPC.

The installation program cannot:

- Subdivide network ranges for the cluster to use.
- Set route tables for the subnets.
- Set VPC options like DHCP.

You must complete these tasks before you install the cluster. See [VPC networking components](#) and [Route tables for your VPC](#) for more information on configuring networking in an AWS VPC.

Your VPC must meet the following characteristics:

- The VPC must not use the **kubernetes.io/cluster/.*: owned, Name**, and **openshift.io/cluster** tags.
The installation program modifies your subnets to add the **kubernetes.io/cluster/.*: shared** tag, so your subnets must have at least one free tag slot available for it. See [Tag Restrictions](#) in the AWS documentation to confirm that the installation program can add a tag to each subnet that you specify. You cannot use a **Name** tag, because it overlaps with the EC2 **Name** field and the installation fails.
- If you want to extend your OpenShift Container Platform cluster into an AWS Outpost and have an existing Outpost subnet, the existing subnet must use the **kubernetes.io/cluster/unmanaged: true** tag. If you do not apply this tag, the installation might fail due to the Cloud Controller Manager creating a service load balancer in the Outpost subnet, which is an unsupported configuration.
- You must enable the **enableDnsSupport** and **enableDnsHostnames** attributes in your VPC, so that the cluster can use the Route 53 zones that are attached to the VPC to resolve cluster's internal DNS records. See [DNS Support in Your VPC](#) in the AWS documentation.
If you prefer to use your own Route 53 hosted private zone, you must associate the existing hosted zone with your VPC prior to installing a cluster. You can define your hosted zone using the **platform.aws.hostedZone** and **platform.aws.hostedZoneRole** fields in the **install-config.yaml** file. You can use a private hosted zone from another account by sharing it with the account where you install the cluster. If you use a private hosted zone from another account, you must use the **Passthrough** or **Manual** credentials mode.

A cluster in an SC2S or C2S Region is unable to reach the public IP addresses for the EC2, ELB, and S3 endpoints. Depending on the level to which you want to restrict internet traffic during the installation, the following configuration options are available:

Option 1: Create VPC endpoints

Create a VPC endpoint and attach it to the subnets that the clusters are using. Name the endpoints as follows:

SC2S

- **elasticloadbalancing.<aws_region>.sc2s.sgov.gov**
- **ec2.<aws_region>.sc2s.sgov.gov**
- **s3.<aws_region>.sc2s.sgov.gov**

C2S

- **elasticloadbalancing.<aws_region>.c2s.ic.gov**
- **ec2.<aws_region>.c2s.ic.gov**
- **s3.<aws_region>.c2s.ic.gov**

With this option, network traffic remains private between your VPC and the required AWS services.

Option 2: Create a proxy without VPC endpoints

As part of the installation process, you can configure an HTTP or HTTPS proxy. With this option, internet traffic goes through the proxy to reach the required AWS services.

Option 3: Create a proxy with VPC endpoints

As part of the installation process, you can configure an HTTP or HTTPS proxy with VPC endpoints. Create a VPC endpoint and attach it to the subnets that the clusters are using. Name the endpoints as follows:

SC2S

- **elasticloadbalancing.<aws_region>.sc2s.sgov.gov**
- **ec2.<aws_region>.sc2s.sgov.gov**
- **s3.<aws_region>.sc2s.sgov.gov**

C2S

- **elasticloadbalancing.<aws_region>.c2s.ic.gov**
- **ec2.<aws_region>.c2s.ic.gov**
- **s3.<aws_region>.c2s.ic.gov**

When configuring the proxy in the **install-config.yaml** file, add these endpoints to the **noProxy** field. With this option, the proxy prevents the cluster from accessing the internet directly. However, network traffic remains private between your VPC and the required AWS services.

Required VPC components

You must provide a suitable VPC and subnets that allow communication to your machines.

Component	AWS type	Description
VPC	<ul style="list-style-type: none"> • AWS::EC2::VPC • AWS::EC2::VPCEndpoint 	You must provide a public VPC for the cluster to use. The VPC uses an endpoint that references the route tables for each subnet to improve communication with the registry that is hosted in S3.

Component	AWS type	Description												
Public subnets	<ul style="list-style-type: none"> ● AWS::EC2::Subnet ● AWS::EC2::SubnetNetworkAclAssociation 	Your VPC must have public subnets for between 1 and 3 availability zones and associate them with appropriate Ingress rules.												
Internet gateway	<ul style="list-style-type: none"> ● AWS::EC2::InternetGateway ● AWS::EC2::VPCGatewayAttachment ● AWS::EC2::RouteTable ● AWS::EC2::Route ● AWS::EC2::SubnetRouteTableAssociation ● AWS::EC2::NatGateway ● AWS::EC2::EIP 	You must have a public internet gateway, with public routes, attached to the VPC. In the provided templates, each public subnet has a NAT gateway with an EIP address. These NAT gateways allow cluster resources, like private subnet instances, to reach the internet and are not required for some restricted network or proxy scenarios.												
Network access control	<ul style="list-style-type: none"> ● AWS::EC2::NetworkAcl ● AWS::EC2::NetworkAclEntry 	<p>You must allow the VPC to access the following ports:</p> <table border="1"> <thead> <tr> <th>Port</th> <th>Reason</th> </tr> </thead> <tbody> <tr> <td>80</td> <td>Inbound HTTP traffic</td> </tr> <tr> <td>443</td> <td>Inbound HTTPS traffic</td> </tr> <tr> <td>22</td> <td>Inbound SSH traffic</td> </tr> <tr> <td>1024 - 65535</td> <td>Inbound ephemeral traffic</td> </tr> <tr> <td>0 - 65535</td> <td>Outbound ephemeral traffic</td> </tr> </tbody> </table>	Port	Reason	80	Inbound HTTP traffic	443	Inbound HTTPS traffic	22	Inbound SSH traffic	1024 - 65535	Inbound ephemeral traffic	0 - 65535	Outbound ephemeral traffic
Port	Reason													
80	Inbound HTTP traffic													
443	Inbound HTTPS traffic													
22	Inbound SSH traffic													
1024 - 65535	Inbound ephemeral traffic													
0 - 65535	Outbound ephemeral traffic													
Private subnets	<ul style="list-style-type: none"> ● AWS::EC2::Subnet ● AWS::EC2::RouteTable ● AWS::EC2::SubnetRouteTableAssociation 	Your VPC can have private subnets. The provided CloudFormation templates can create private subnets for between 1 and 3 availability zones. If you use private subnets, you must provide appropriate routes and tables for them.												

6.3.9.5.2. VPC validation

To ensure that the subnets that you provide are suitable, the installation program confirms the following data:

- All the subnets that you specify exist.
- You provide private subnets.
- The subnet CIDRs belong to the machine CIDR that you specified.
- You provide subnets for each availability zone. Each availability zone contains no more than one public and one private subnet. If you use a private cluster, provide only a private subnet for each availability zone. Otherwise, provide exactly one public and private subnet for each availability zone.
- You provide a public subnet for each private subnet availability zone. Machines are not provisioned in availability zones that you do not provide private subnets for.

If you destroy a cluster that uses an existing VPC, the VPC is not deleted. When you remove the OpenShift Container Platform cluster from a VPC, the **kubernetes.io/cluster/.*: shared** tag is removed from the subnets that it used.

6.3.9.5.3. Division of permissions

Starting with OpenShift Container Platform 4.3, you do not need all of the permissions that are required for an installation program-provisioned infrastructure cluster to deploy a cluster. This change mimics the division of permissions that you might have at your company: some individuals can create different resource in your clouds than others. For example, you might be able to create application-specific items, like instances, buckets, and load balancers, but not networking-related components such as VPCs, subnets, or ingress rules.

The AWS credentials that you use when you create your cluster do not need the networking permissions that are required to make VPCs and core networking components within the VPC, such as subnets, routing tables, internet gateways, NAT, and VPN. You still need permission to make the application resources that the machines within the cluster require, such as ELBs, security groups, S3 buckets, and nodes.

6.3.9.5.4. Isolation between clusters

If you deploy OpenShift Container Platform to an existing network, the isolation of cluster services is reduced in the following ways:

- You can install multiple OpenShift Container Platform clusters in the same VPC.
- ICMP ingress is allowed from the entire network.
- TCP 22 ingress (SSH) is allowed to the entire network.
- Control plane TCP 6443 ingress (Kubernetes API) is allowed to the entire network.
- Control plane TCP 22623 ingress (MCS) is allowed to the entire network.

6.3.9.5.5. Optional: AWS security groups

By default, the installation program creates and attaches security groups to control plane and compute machines. The rules associated with the default security groups cannot be modified.

However, you can apply additional existing AWS security groups, which are associated with your existing VPC, to control plane and compute machines. Applying custom security groups can help you meet the security needs of your organization, in such cases where you need to control the incoming or outgoing traffic of these machines.

As part of the installation process, you apply custom security groups by modifying the **install-config.yaml** file before deploying the cluster.

For more information, see "Applying existing AWS security groups to the cluster".

6.3.9.6. Uploading a custom RHCOS AMI in AWS

If you are deploying to a custom Amazon Web Services (AWS) region, you must upload a custom Red Hat Enterprise Linux CoreOS (RHCOS) Amazon Machine Image (AMI) that belongs to that region.

Prerequisites

- You configured an AWS account.
- You created an Amazon S3 bucket with the required IAM [service role](#).
- You uploaded your RHCOS VMDK file to Amazon S3. The RHCOS VMDK file must be the highest version that is less than or equal to the OpenShift Container Platform version you are installing.
- You downloaded the AWS CLI and installed it on your computer. See [Install the AWS CLI Using the Bundled Installer](#).

Procedure

1. Export your AWS profile as an environment variable:

```
$ export AWS_PROFILE=<aws_profile> 1
```

2. Export the region to associate with your custom AMI as an environment variable:

```
$ export AWS_DEFAULT_REGION=<aws_region> 1
```

3. Export the version of RHCOS you uploaded to Amazon S3 as an environment variable:

```
$ export RHCOS_VERSION=<version> 1
```

1 1 1 The RHCOS VMDK version, like **4.16.0**.

4. Export the Amazon S3 bucket name as an environment variable:

```
$ export VMIMPORT_BUCKET_NAME=<s3_bucket_name>
```

5. Create the **containers.json** file and define your RHCOS VMDK file:

```
$ cat <<EOF > containers.json
{
  "Description": "rhcos-${RHCOS_VERSION}-x86_64-aws.x86_64",
```



```
"Format": "vmdk",
"UserBucket": {
  "S3Bucket": "${VMIMPORT_BUCKET_NAME}",
  "S3Key": "rhcos-${RHCOS_VERSION}-x86_64-aws.x86_64.vmdk"
}
}
EOF
```

6. Import the RHCOS disk as an Amazon EBS snapshot:

```
$ aws ec2 import-snapshot --region ${AWS_DEFAULT_REGION} \
  --description "<description>" \ 1
  --disk-container "file://<file_path>/containers.json" 2
```

1 The description of your RHCOS disk being imported, like **rhcos-\${RHCOS_VERSION}-x86_64-aws.x86_64**.

2 The file path to the JSON file describing your RHCOS disk. The JSON file should contain your Amazon S3 bucket name and key.

7. Check the status of the image import:

```
$ watch -n 5 aws ec2 describe-import-snapshot-tasks --region ${AWS_DEFAULT_REGION}
```

Example output

```
{
  "ImportSnapshotTasks": [
    {
      "Description": "rhcos-4.7.0-x86_64-aws.x86_64",
      "ImportTaskId": "import-snap-fh6i8uil",
      "SnapshotTaskDetail": {
        "Description": "rhcos-4.7.0-x86_64-aws.x86_64",
        "DiskImageSize": 819056640.0,
        "Format": "VMDK",
        "SnapshotId": "snap-06331325870076318",
        "Status": "completed",
        "UserBucket": {
          "S3Bucket": "external-images",
          "S3Key": "rhcos-4.7.0-x86_64-aws.x86_64.vmdk"
        }
      }
    }
  ]
}
```

Copy the **SnapshotId** to register the image.

8. Create a custom RHCOS AMI from the RHCOS snapshot:

```
$ aws ec2 register-image \
  --region ${AWS_DEFAULT_REGION} \
  --architecture x86_64 \ 1
```

```
--description "rhcos-${RHCOS_VERSION}-x86_64-aws.x86_64" \ 2
--ena-support \
--name "rhcos-${RHCOS_VERSION}-x86_64-aws.x86_64" \ 3
--virtualization-type hvm \
--root-device-name '/dev/xvda' \
--block-device-mappings 'DeviceName=/dev/xvda,Ebs=
{DeleteOnTermination=true,SnapshotId=<snapshot_ID>}' 4
```

- 1 The RHCOS VMDK architecture type, like **x86_64**, **aarch64**, **s390x**, or **ppc64le**.
- 2 The **Description** from the imported snapshot.
- 3 The name of the RHCOS AMI.
- 4 The **SnapshotID** from the imported snapshot.

To learn more about these APIs, see the AWS documentation for [importing snapshots](#) and [creating EBS-backed AMIs](#).

6.3.9.7. Manually creating the installation configuration file

Installing the cluster requires that you manually create the installation configuration file.

Prerequisites

- You have uploaded a custom RHCOS AMI.
- You have an SSH public key on your local machine to provide to the installation program. The key will be used for SSH authentication onto your cluster nodes for debugging and disaster recovery.
- You have obtained the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create an installation directory to store your required installation assets in:

```
$ mkdir <installation_directory>
```



IMPORTANT


You must create a directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

2. Customize the sample **install-config.yaml** file template that is provided and save it in the **<installation_directory>**.

**NOTE**

You must name this configuration file **install-config.yaml**.

3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.

**IMPORTANT**


The **install-config.yaml** file is consumed during the next step of the installation process. You must back it up now.

Additional resources

- [Installation configuration parameters for AWS](#)

6.3.9.7.1. Tested instance types for AWS

The following Amazon Web Services (AWS) instance types have been tested with OpenShift Container Platform.

**NOTE**


Use the machine types included in the following charts for your AWS instances. If you use an instance type that is not listed in the chart, ensure that the instance size you use matches the minimum resource requirements that are listed in the section named "Minimum resource requirements for cluster installation".

Example 6.41. Machine types based on 64-bit x86 architecture for select regions

- **c4.***
- **c5.***
- **i3.***
- **m4.***
- **m5.***
- **r4.***
- **r5.***
- **t3.***

6.3.9.7.2. Sample customized install-config.yaml file for AWS

You can customize the installation configuration file (**install-config.yaml**) to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.



IMPORTANT

This sample YAML file is provided for reference only. Use it as a resource to enter parameter values into the installation configuration file that you created manually.

```

apiVersion: v1
baseDomain: example.com 1
credentialsMode: Mint 2
controlPlane: 3 4
  hyperthreading: Enabled 5
  name: master
  platform:
    aws:
      zones:
        - us-iso-east-1a
        - us-iso-east-1b
      rootVolume:
        iops: 4000
        size: 500
        type: io1 6
      metadataService:
        authentication: Optional 7
        type: m6i.xlarge
      replicas: 3
compute: 8
- hyperthreading: Enabled 9
  name: worker
  platform:
    aws:
      rootVolume:
        iops: 2000
        size: 500
        type: io1 10
      metadataService:
        authentication: Optional 11
        type: c5.4xlarge
      zones:
        - us-iso-east-1a
        - us-iso-east-1b
      replicas: 3
  metadata:
    name: test-cluster 12
  networking:
    clusterNetwork:
      - cidr: 10.128.0.0/14
      hostPrefix: 23
    machineNetwork:
      - cidr: 10.0.0.0/16
    networkType: OVNKubernetes 13
    serviceNetwork:
      - 172.30.0.0/16
  platform:
    aws:
      region: us-iso-east-1 14

```

```

propagateUserTags: true 15
userTags:
  adminContact: jdoe
  costCenter: 7536
subnets: 16
- subnet-1
- subnet-2
- subnet-3
amiID: ami-96c6f8f7 17 18
serviceEndpoints: 19
- name: ec2
  url: https://vpce-id.ec2.us-west-2.vpce.amazonaws.com
hostedZone: Z3URY6TWQ91KVV 20
fips: false 21
sshKey: ssh-ed25519 AAAA... 22
publish: Internal 23
pullSecret: '{"auths": ...}' 24
additionalTrustBundle: | 25
-----BEGIN CERTIFICATE-----
<MY_TRUSTED_CA_CERT>
-----END CERTIFICATE-----

```

1 12 14 17 24 Required.

2 Optional: Add this parameter to force the Cloud Credential Operator (CCO) to use the specified mode. By default, the CCO uses the root credentials in the **kube-system** namespace to dynamically try to determine the capabilities of the credentials. For details about CCO modes, see the "About the Cloud Credential Operator" section in the *Authentication and authorization* guide.

3 8 15 If you do not provide these parameters and values, the installation program provides the default value.

4 The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.

5 9 Whether to enable or disable simultaneous multithreading, or **hyperthreading**. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores. You can disable it by setting the parameter value to **Disabled**. If you disable simultaneous multithreading in some cluster machines, you must disable it in all cluster machines.



IMPORTANT

If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance. Use larger instance types, such as **m4.2xlarge** or **m5.2xlarge**, for your machines if you disable simultaneous multithreading.

6 10 To configure faster storage for etcd, especially for larger clusters, set the storage type as **io1** and set **iops** to **2000**.

7 11

Whether to require the [Amazon EC2 Instance Metadata Service v2](#) (IMDSv2). To require IMDSv2, set the parameter value to **Required**. To allow the use of both IMDSv1 and IMDSv2, set the



NOTE

The IMDS configuration for control plane machines that is set during cluster installation can only be changed by using the AWS CLI. The IMDS configuration for compute machines can be changed by using compute machine sets.

- 13 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 16 If you provide your own VPC, specify subnets for each availability zone that your cluster uses.
- 18 The ID of the AMI used to boot machines for the cluster. If set, the AMI must belong to the same region as the cluster.
- 19 The AWS service endpoints. Custom endpoints are required when installing to an unknown AWS region. The endpoint URL must use the **https** protocol and the host must trust the certificate.
- 20 The ID of your existing Route 53 private hosted zone. Providing an existing hosted zone requires that you supply your own VPC and the hosted zone is already associated with the VPC prior to installing your cluster. If undefined, the installation program creates a new hosted zone.
- 21 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.



IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 22 You can optionally provide the **sshKey** value that you use to access the machines in your cluster.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- 23 How to publish the user-facing endpoints of your cluster. Set **publish** to **Internal** to deploy a private cluster, which cannot be accessed from the internet. The default value is **External**.

- 25

The custom CA certificate. This is required when deploying to the SC2S or C2S Regions because the AWS API requires a custom CA trust bundle.

6.3.9.7.3. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: ec2.<aws_region>.amazonaws.com,elasticloadbalancing.
<aws_region>.amazonaws.com,s3.<aws_region>.amazonaws.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all

destinations. If you have added the Amazon **EC2**, **Elastic Load Balancing**, and **S3** VPC endpoints to your VPC, you must add these endpoints to the **noProxy** field.

- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.



NOTE

The installation program does not support the proxy **readinessEndpoints** field.



NOTE

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.



NOTE

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

6.3.9.7.4. Applying existing AWS security groups to the cluster

Applying existing AWS security groups to your control plane and compute machines can help you meet the security needs of your organization, in such cases where you need to control the incoming or outgoing traffic of these machines.

Prerequisites

- You have created the security groups in AWS. For more information, see the AWS documentation about working with [security groups](#).
- The security groups must be associated with the existing VPC that you are deploying the cluster to. The security groups cannot be associated with another VPC.

- You have an existing **install-config.yaml** file.

Procedure

1. In the **install-config.yaml** file, edit the **compute.platform.aws.additionalSecurityGroupIDs** parameter to specify one or more custom security groups for your compute machines.
2. Edit the **controlPlane.platform.aws.additionalSecurityGroupIDs** parameter to specify one or more custom security groups for your control plane machines.
3. Save the file and reference it when deploying the cluster.

Sample **install-config.yaml** file that specifies custom security groups

```
# ...
compute:
- hyperthreading: Enabled
  name: worker
  platform:
    aws:
      additionalSecurityGroupIDs:
        - sg-1 1
        - sg-2
  replicas: 3
controlPlane:
  hyperthreading: Enabled
  name: master
  platform:
    aws:
      additionalSecurityGroupIDs:
        - sg-3
        - sg-4
  replicas: 3
platform:
  aws:
    region: us-east-1
  subnets: 2
    - subnet-1
    - subnet-2
    - subnet-3
```

- 1** Specify the name of the security group as it appears in the Amazon EC2 console, including the **sg** prefix.
- 2** Specify subnets for each availability zone that your cluster uses.

6.3.9.8. Alternatives to storing administrator-level secrets in the kube-system project

By default, administrator secrets are stored in the **kube-system** project. If you configured the **credentialsMode** parameter in the **install-config.yaml** file to **Manual**, you must use one of the following alternatives:

- To manage long-term cloud credentials manually, follow the procedure in [Manually creating long-term credentials](#).

- To implement short-term credentials that are managed outside the cluster for individual components, follow the procedures in [Configuring an AWS cluster to use short-term credentials](#).

6.3.9.8.1. Manually creating long-term credentials

The Cloud Credential Operator (CCO) can be put into manual mode prior to installation in environments where the cloud identity and access management (IAM) APIs are not reachable, or the administrator prefers not to store an administrator-level credential secret in the cluster **kube-system** namespace.

Procedure

1. If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```
apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...
```

2. If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

3. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

4. Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included \ ❶
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \ ❷
  --to=<path_to_directory_for_credentials_requests> \ ❸
```

- ❶ The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- ❷ Specify the location of the **install-config.yaml** file.
- ❸ Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

This command creates a YAML file for each **CredentialsRequest** object.

Sample **CredentialsRequest** object

```

apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: AWSProviderSpec
    statementEntries:
      - effect: Allow
        action:
          - iam:GetUser
          - iam:GetUserPolicy
          - iam:ListAccessKeys
        resource: "*"
    ...
  ...

```

5. Create YAML files for secrets in the **openshift-install** manifests directory that you generated previously. The secrets must be stored using the namespace and secret name defined in the **spec.secretRef** for each **CredentialsRequest** object.

Sample **CredentialsRequest** object with secrets

```

apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: AWSProviderSpec
    statementEntries:
      - effect: Allow
        action:
          - s3:CreateBucket
          - s3>DeleteBucket
        resource: "*"
    ...
  secretRef:
    name: <component_secret>
    namespace: <component_namespace>
  ...

```

Sample **Secret** object

```


apiVersion: v1

```

```

kind: Secret
metadata:
  name: <component_secret>
  namespace: <component_namespace>
data:
  aws_access_key_id: <base64_encoded_aws_access_key_id>
  aws_secret_access_key: <base64_encoded_aws_secret_access_key>

```



IMPORTANT


Before upgrading a cluster that uses manually maintained credentials, you must ensure that the CCO is in an upgradeable state.

6.3.9.8.2. Configuring an AWS cluster to use short-term credentials

To install a cluster that is configured to use the AWS Security Token Service (STS), you must configure the CCO utility and create the required AWS resources for your cluster.

6.3.9.8.2.1. Configuring the Cloud Credential Operator utility

To create and manage cloud credentials from outside of the cluster when the Cloud Credential Operator (CCO) is operating in manual mode, extract and prepare the CCO utility (**ccoctl**) binary.



NOTE

The **ccoctl** utility is a Linux binary that must run in a Linux environment.

Prerequisites

- You have access to an OpenShift Container Platform account with cluster administrator access.
- You have installed the OpenShift CLI (**oc**).
- You have created an AWS account for the **ccoctl** utility to use with the following permissions:

Example 6.42. Required AWS permissions

Required iam permissions

- **iam:CreateOpenIDConnectProvider**
- **iam:CreateRole**
- **iam>DeleteOpenIDConnectProvider**
- **iam>DeleteRole**
- **iam>DeleteRolePolicy**
- **iam:GetOpenIDConnectProvider**
- **iam:GetRole**
- **iam:GetUser**
- **iam:ListOpenIDConnectProviders**

- **iam:ListRolePolicies**
- **iam:ListRoles**
- **iam:PutRolePolicy**
- **iam:TagOpenIDConnectProvider**
- **iam:TagRole**

Required s3 permissions

- **s3:CreateBucket**
- **s3>DeleteBucket**
- **s3>DeleteObject**
- **s3:GetBucketAcl**
- **s3:GetBucketTagging**
- **s3:GetObject**
- **s3:GetObjectAcl**
- **s3:GetObjectTagging**
- **s3:ListBucket**
- **s3:PutBucketAcl**
- **s3:PutBucketPolicy**
- **s3:PutBucketPublicAccessBlock**
- **s3:PutBucketTagging**
- **s3:PutObject**
- **s3:PutObjectAcl**
- **s3:PutObjectTagging**

Required cloudfront permissions

- **cloudfront:ListCloudFrontOriginAccessIdentities**
- **cloudfront:ListDistributions**
- **cloudfront:ListTagsForResource**

If you plan to store the OIDC configuration in a private S3 bucket that is accessed by the IAM identity provider through a public CloudFront distribution URL, the AWS account that runs the **ccoctl** utility requires the following additional permissions:

Example 6.43. Additional permissions for a private S3 bucket with CloudFront

- **cloudfront:CreateCloudFrontOriginAccessIdentity**
- **cloudfront:CreateDistribution**
- **cloudfront>DeleteCloudFrontOriginAccessIdentity**
- **cloudfront>DeleteDistribution**
- **cloudfront:GetCloudFrontOriginAccessIdentity**
- **cloudfront:GetCloudFrontOriginAccessIdentityConfig**
- **cloudfront:GetDistribution**
- **cloudfront:TagResource**
- **cloudfront:UpdateDistribution**

**NOTE**

These additional permissions support the use of the **--create-private-s3-bucket** option when processing credentials requests with the **ccoctl aws create-all** command.


Procedure

1. Set a variable for the OpenShift Container Platform release image by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

2. Obtain the CCO container image from the OpenShift Container Platform release image by running the following command:

```
$ CCO_IMAGE=$(oc adm release info --image-for='cloud-credential-operator'
$RELEASE_IMAGE -a ~/.pull-secret)
```

**NOTE**

Ensure that the architecture of the **\$RELEASE_IMAGE** matches the architecture of the environment in which you will use the **ccoctl** tool.

3. Extract the **ccoctl** binary from the CCO container image within the OpenShift Container Platform release image by running the following command:

```
$ oc image extract $CCO_IMAGE \
--file="/usr/bin/ccoctl.<rhel_version>" \
-a ~/.pull-secret
```

- 1 For **<rhel_version>**, specify the value that corresponds to the version of Red Hat Enterprise Linux (RHEL) that the host uses. If no value is specified, **ccoctl.rhel8** is used by default. The following values are valid:
 - **rhel8**: Specify this value for hosts that use RHEL 8.
 - **rhel9**: Specify this value for hosts that use RHEL 9.

4. Change the permissions to make **ccoctl** executable by running the following command:

```
$ chmod 775 ccoctl.<rhel_version>
```

Verification

- To verify that **ccoctl** is ready to use, display the help file by running the following command:

```
$ ccoctl --help
```

Output of **ccoctl --help**

```
OpenShift credentials provisioning tool
```

```
Usage:
```

```
ccoctl [command]
```

```
Available Commands:
```

```
aws      Manage credentials objects for AWS cloud
azure    Manage credentials objects for Azure
gcp      Manage credentials objects for Google cloud
help     Help about any command
ibmcloud Manage credentials objects for IBM Cloud
nutanix  Manage credentials objects for Nutanix
```

```
Flags:
```

```
-h, --help  help for ccoctl
```

```
Use "ccoctl [command] --help" for more information about a command.
```

6.3.9.8.2.2. Creating AWS resources with the Cloud Credential Operator utility

You have the following options when creating AWS resources:

- You can use the **ccoctl aws create-all** command to create the AWS resources automatically. This is the quickest way to create the resources. See [Creating AWS resources with a single command](#).
- If you need to review the JSON files that the **ccoctl** tool creates before modifying AWS resources, or if the process the **ccoctl** tool uses to create AWS resources automatically does not meet the requirements of your organization, you can create the AWS resources individually. See [Creating AWS resources individually](#).

6.3.9.8.2.2.1. Creating AWS resources with a single command

If the process the **ccoctl** tool uses to create AWS resources automatically meets the requirements of your organization, you can use the **ccoctl aws create-all** command to automate the creation of AWS resources.

Otherwise, you can create the AWS resources individually. For more information, see "Creating AWS resources individually".



NOTE

By default, **ccoctl** creates objects in the directory in which the commands are run. To create the objects in a different directory, use the **--output-dir** flag. This procedure uses **<path_to_ccoctl_output_dir>** to refer to this directory.

Prerequisites

You must have:

- Extracted and prepared the **ccoctl** binary.

Procedure

1. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

2. Extract the list of **CredentialsRequest** objects from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included \1
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \2
  --to=<path_to_directory_for_credentials_requests> \3
```

- 1 The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2 Specify the location of the **install-config.yaml** file.
- 3 Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.



NOTE

This command might take a few moments to run.

3. Use the **ccoctl** tool to process all **CredentialsRequest** objects by running the following command:

```
$ ccoctl aws create-all \
  --name=<name> \1
```



```

--region=<aws_region> \ 2
--credentials-requests-dir=<path_to_credentials_requests_directory> \ 3
--output-dir=<path_to_ccoctl_output_dir> \ 4
--create-private-s3-bucket 5

```

- 1 Specify the name used to tag any cloud resources that are created for tracking.
- 2 Specify the AWS region in which cloud resources will be created.
- 3 Specify the directory containing the files for the component **CredentialsRequest** objects.
- 4 Optional: Specify the directory in which you want the **ccoctl** utility to create objects. By default, the utility creates objects in the directory in which the commands are run.
- 5 Optional: By default, the **ccoctl** utility stores the OpenID Connect (OIDC) configuration files in a public S3 bucket and uses the S3 URL as the public OIDC endpoint. To store the OIDC configuration in a private S3 bucket that is accessed by the IAM identity provider through a public CloudFront distribution URL instead, use the **--create-private-s3-bucket** parameter.



NOTE

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

Verification

- To verify that the OpenShift Container Platform secrets are created, list the files in the **<path_to_ccoctl_output_dir>/manifests** directory:

```
$ ls <path_to_ccoctl_output_dir>/manifests
```

Example output

```

cluster-authentication-02-config.yaml
openshift-cloud-credential-operator-cloud-credential-operator-iam-ro-creds-credentials.yaml
openshift-cloud-network-config-controller-cloud-credentials-credentials.yaml
openshift-cluster-api-capamanager-bootstrap-credentials-credentials.yaml
openshift-cluster-csi-drivers-ebs-cloud-credentials-credentials.yaml
openshift-image-registry-installer-cloud-credentials-credentials.yaml
openshift-ingress-operator-cloud-credentials-credentials.yaml
openshift-machine-api-aws-cloud-credentials-credentials.yaml

```

You can verify that the IAM roles are created by querying AWS. For more information, refer to AWS documentation on listing IAM roles.

6.3.9.8.2.2.2. Creating AWS resources individually

You can use the **ccoctl** tool to create AWS resources individually. This option might be useful for an organization that shares the responsibility for creating these resources among different users or departments.

Otherwise, you can use the **ccoctl aws create-all** command to create the AWS resources automatically. For more information, see "Creating AWS resources with a single command".



NOTE

By default, **ccoctl** creates objects in the directory in which the commands are run. To create the objects in a different directory, use the **--output-dir** flag. This procedure uses **<path_to_ccoctl_output_dir>** to refer to this directory.

Some **ccoctl** commands make AWS API calls to create or modify AWS resources. You can use the **--dry-run** flag to avoid making API calls. Using this flag creates JSON files on the local file system instead. You can review and modify the JSON files and then apply them with the AWS CLI tool using the **--cli-input-json** parameters.

Prerequisites

- Extract and prepare the **ccoctl** binary.

Procedure

1. Generate the public and private RSA key files that are used to set up the OpenID Connect provider for the cluster by running the following command:

```
$ ccoctl aws create-key-pair
```

Example output

```
2021/04/13 11:01:02 Generating RSA keypair
2021/04/13 11:01:03 Writing private key to /<path_to_ccoctl_output_dir>/serviceaccount-signer.private
2021/04/13 11:01:03 Writing public key to /<path_to_ccoctl_output_dir>/serviceaccount-signer.public
2021/04/13 11:01:03 Copying signing key for use by installer
```

where **serviceaccount-signer.private** and **serviceaccount-signer.public** are the generated key files.

This command also creates a private key that the cluster requires during installation in **/<path_to_ccoctl_output_dir>/tls/bound-service-account-signing-key.key**.

2. Create an OpenID Connect identity provider and S3 bucket on AWS by running the following command:

```
$ ccoctl aws create-identity-provider \
  --name=<name> \ 1
  --region=<aws_region> \ 2
  --public-key-file=<path_to_ccoctl_output_dir>/serviceaccount-signer.public 3
```

- 1** **<name>** is the name used to tag any cloud resources that are created for tracking.
- 2** **<aws-region>** is the AWS region in which cloud resources will be created.
- 3** **<path_to_ccoctl_output_dir>** is the path to the public key file that the **ccoctl aws create-key-pair** command generated.

Example output

```
2021/04/13 11:16:09 Bucket <name>-oidc created
2021/04/13 11:16:10 OpenID Connect discovery document in the S3 bucket <name>-oidc at
.well-known/openid-configuration updated
2021/04/13 11:16:10 Reading public key
2021/04/13 11:16:10 JSON web key set (JWKS) in the S3 bucket <name>-oidc at keys.json
updated
2021/04/13 11:16:18 Identity Provider created with ARN: arn:aws:iam::
<aws_account_id>:oidc-provider/<name>-oidc.s3.<aws_region>.amazonaws.com
```

where **openid-configuration** is a discovery document and **keys.json** is a JSON web key set file.

This command also creates a YAML configuration file in `/<path_to_ccoctl_output_dir>/manifests/cluster-authentication-02-config.yaml`. This file sets the issuer URL field for the service account tokens that the cluster generates, so that the AWS IAM identity provider trusts the tokens.

3. Create IAM roles for each component in the cluster:
 - a. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

- b. Extract the list of **CredentialsRequest** objects from the OpenShift Container Platform release image:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included 1 \
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \
2 \
  --to=<path_to_directory_for_credentials_requests> 3
```

- 1** The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2** Specify the location of the **install-config.yaml** file.
- 3** Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

- c. Use the **ccoctl** tool to process all **CredentialsRequest** objects by running the following command:

```
$ ccoctl aws create-iam-roles \
  --name=<name> \
  --region=<aws_region> \
  --credentials-requests-dir=<path_to_credentials_requests_directory> \
  --identity-provider-arn=arn:aws:iam::<aws_account_id>:oidc-provider/<name>-oidc.s3.
<aws_region>.amazonaws.com
```



NOTE

For AWS environments that use alternative IAM API endpoints, such as GovCloud, you must also specify your region with the **--region** parameter.

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

For each **CredentialsRequest** object, **ccoctl** creates an IAM role with a trust policy that is tied to the specified OIDC identity provider, and a permissions policy as defined in each **CredentialsRequest** object from the OpenShift Container Platform release image.

Verification

- To verify that the OpenShift Container Platform secrets are created, list the files in the **<path_to_ccoctl_output_dir>/manifests** directory:

```
$ ls <path_to_ccoctl_output_dir>/manifests
```

Example output

```
cluster-authentication-02-config.yaml
openshift-cloud-credential-operator-cloud-credential-operator-iam-ro-creds-credentials.yaml
openshift-cloud-network-config-controller-cloud-credentials-credentials.yaml
openshift-cluster-api-capamanager-bootstrap-credentials-credentials.yaml
openshift-cluster-csi-drivers-ebs-cloud-credentials-credentials.yaml
openshift-image-registry-installer-cloud-credentials-credentials.yaml
openshift-ingress-operator-cloud-credentials-credentials.yaml
openshift-machine-api-aws-cloud-credentials-credentials.yaml
```

You can verify that the IAM roles are created by querying AWS. For more information, refer to AWS documentation on listing IAM roles.

6.3.9.8.2.3. Incorporating the Cloud Credential Operator utility manifests

To implement short-term security credentials managed outside the cluster for individual components, you must move the manifest files that the Cloud Credential Operator utility (**ccoctl**) created to the correct directories for the installation program.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have configured the Cloud Credential Operator utility (**ccoctl**).
- You have created the cloud provider resources that are required for your cluster with the **ccoctl** utility.

Procedure

- If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```
apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...
```

- If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

- Copy the manifests that the **ccoctl** utility generated to the **manifests** directory that the installation program created by running the following command:

```
$ cp /<path_to_ccoctl_output_dir>/manifests/* ./manifests/
```

- Copy the private key that the **ccoctl** utility generated in the **tls** directory to the installation directory by running the following command:

```
$ cp -a /<path_to_ccoctl_output_dir>/tls .
```

6.3.9.9. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.

Procedure

- Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

- 1 For `<installation_directory>`, specify the location of your customized `./install-config.yaml` file.
 - 2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.
2. Optional: Remove or disable the **AdministratorAccess** policy from the IAM account that you used to install the cluster.




NOTE

The elevated permissions provided by the **AdministratorAccess** policy are required only during installation.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to `<installation_directory>/openshift_install.log`.




IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

6.3.9.10. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

6.3.9.11. Logging in to the cluster by using the web console

The **kubeadmin** user exists by default after an OpenShift Container Platform installation. You can log in to your cluster as the **kubeadmin** user by using the OpenShift Container Platform web console.


Prerequisites

- You have access to the installation host.
- You completed a cluster installation and all cluster Operators are available.

Procedure

1. Obtain the password for the **kubeadmin** user from the **kubeadmin-password** file on the installation host:

```
$ cat <installation_directory>/auth/kubeadmin-password
```



NOTE

Alternatively, you can obtain the **kubeadmin** password from the **<installation_directory>/openshift_install.log** log file on the installation host.

2. List the OpenShift Container Platform web console route:

```
$ oc get routes -n openshift-console | grep 'console-openshift'
```



NOTE

Alternatively, you can obtain the OpenShift Container Platform route from the `<installation_directory>/openshift_install.log` log file on the installation host.

Example output

```
console    console-openshift-console.apps.<cluster_name>.<base_domain>    console
https    reencrypt/Redirect    None
```

3. Navigate to the route detailed in the output of the preceding command in a web browser and log in as the **kubeadmin** user.

Additional resources

- [Accessing the web console](#)

6.3.9.12. Next steps

- [Validating an installation.](#)
- [Customize your cluster.](#)
- If necessary, you can [opt out of remote health reporting](#) .
- If necessary, you can [remove cloud provider credentials](#).

6.3.10. Installing a cluster on AWS China

In OpenShift Container Platform version 4.16, you can install a cluster to the following Amazon Web Services (AWS) China regions:

- **cn-north-1** (Beijing)
- **cn-northwest-1** (Ningxia)

6.3.10.1. Prerequisites

- You have an Internet Content Provider (ICP) license.
- You reviewed details about the [OpenShift Container Platform installation and update processes](#).
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [configured an AWS account](#) to host the cluster.
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.



IMPORTANT

If you have an AWS profile stored on your computer, it must not use a temporary session token that you generated while using a multi-factor authentication device. The cluster continues to use your current AWS credentials to create AWS resources for the entire life of the cluster, so you must use long-term credentials. To generate appropriate keys, see [Managing Access Keys for IAM Users](#) in the AWS documentation. You can supply the keys when you run the installation program.

6.3.10.2. Installation requirements

Red Hat does not publish a Red Hat Enterprise Linux CoreOS (RHCOS) Amazon Machine Image (AMI) for the AWS China regions.

Before you can install the cluster, you must:

- Upload a custom RHCOS AMI.
- Manually create the installation configuration file (**install-config.yaml**).
- Specify the AWS region, and the accompanying custom AMI, in the installation configuration file.

You cannot use the OpenShift Container Platform installation program to create the installation configuration file. The installer does not list an AWS region without native support for an RHCOS AMI.

6.3.10.3. Private clusters

You can deploy a private OpenShift Container Platform cluster that does not expose external endpoints. Private clusters are accessible from only an internal network and are not visible to the internet.

By default, OpenShift Container Platform is provisioned to use publicly-accessible DNS and endpoints. A private cluster sets the DNS, Ingress Controller, and API server to private when you deploy your cluster. This means that the cluster resources are only accessible from your internal network and are not visible to the internet.




IMPORTANT

If the cluster has any public subnets, load balancer services created by administrators might be publicly accessible. To ensure cluster security, verify that these services are explicitly annotated as private.

To deploy a private cluster, you must:

- Use existing networking that meets your requirements. Your cluster resources might be shared between other clusters on the network.
- Deploy from a machine that has access to:
 - The API services for the cloud to which you provision.
 - The hosts on the network that you provision.
 - The internet to obtain installation media.

You can use any machine that meets these access requirements and follows your company's guidelines. For example, this machine can be a bastion host on your cloud network.



NOTE

AWS China does not support a VPN connection between the VPC and your network. For more information about the Amazon VPC service in the Beijing and Ningxia regions, see [Amazon Virtual Private Cloud](#) in the AWS China documentation.

6.3.10.3.1. Private clusters in AWS

To create a private cluster on Amazon Web Services (AWS), you must provide an existing private VPC and subnets to host the cluster. The installation program must also be able to resolve the DNS records that the cluster requires. The installation program configures the Ingress Operator and API server for access from only the private network.

The cluster still requires access to internet to access the AWS APIs.

The following items are not required or created when you install a private cluster:

- Public subnets
- Public load balancers, which support public ingress
- A public Route 53 zone that matches the **baseDomain** for the cluster

The installation program does use the **baseDomain** that you specify to create a private Route 53 zone and the required records for the cluster. The cluster is configured so that the Operators do not create public records for the cluster and all cluster machines are placed in the private subnets that you specify.

6.3.10.3.1.1. Limitations

The ability to add public functionality to a private cluster is limited.

- You cannot make the Kubernetes API endpoints public after installation without taking additional actions, including creating public subnets in the VPC for each availability zone in use, creating a public load balancer, and configuring the control plane security groups to allow traffic from the internet on 6443 (Kubernetes API port).
- If you use a public Service type load balancer, you must tag a public subnet in each availability zone with **kubernetes.io/cluster/<cluster-infra-id>: shared** so that AWS can use them to create public load balancers.

6.3.10.4. About using a custom VPC

In OpenShift Container Platform 4.16, you can deploy a cluster into existing subnets in an existing Amazon Virtual Private Cloud (VPC) in Amazon Web Services (AWS). By deploying OpenShift Container Platform into an existing AWS VPC, you might be able to avoid limit constraints in new accounts or more easily abide by the operational constraints that your company's guidelines set. If you cannot obtain the infrastructure creation permissions that are required to create the VPC yourself, use this installation option.

Because the installation program cannot know what other components are also in your existing subnets, it cannot choose subnet CIDRs and so forth on your behalf. You must configure networking for the subnets that you install your cluster to yourself.

6.3.10.4.1. Requirements for using your VPC

The installation program no longer creates the following components:

- Internet gateways
- NAT gateways
- Subnets
- Route tables
- VPCs
- VPC DHCP options
- VPC endpoints



NOTE

The installation program requires that you use the cloud-provided DNS server. Using a custom DNS server is not supported and causes the installation to fail.

If you use a custom VPC, you must correctly configure it and its subnets for the installation program and the cluster to use. See [Amazon VPC console wizard configurations](#) and [Work with VPCs and subnets](#) in the AWS documentation for more information on creating and managing an AWS VPC.

The installation program cannot:

- Subdivide network ranges for the cluster to use.
- Set route tables for the subnets.
- Set VPC options like DHCP.

You must complete these tasks before you install the cluster. See [VPC networking components](#) and [Route tables for your VPC](#) for more information on configuring networking in an AWS VPC.

Your VPC must meet the following characteristics:

- The VPC must not use the **kubernetes.io/cluster/.*: owned, Name**, and **openshift.io/cluster** tags.
The installation program modifies your subnets to add the **kubernetes.io/cluster/.*: shared** tag, so your subnets must have at least one free tag slot available for it. See [Tag Restrictions](#) in the AWS documentation to confirm that the installation program can add a tag to each subnet that you specify. You cannot use a **Name** tag, because it overlaps with the EC2 **Name** field and the installation fails.
- If you want to extend your OpenShift Container Platform cluster into an AWS Outpost and have an existing Outpost subnet, the existing subnet must use the **kubernetes.io/cluster/unmanaged: true** tag. If you do not apply this tag, the installation might fail due to the Cloud Controller Manager creating a service load balancer in the Outpost subnet, which is an unsupported configuration.
- You must enable the **enableDnsSupport** and **enableDnsHostnames** attributes in your VPC, so that the cluster can use the Route 53 zones that are attached to the VPC to resolve cluster's internal DNS records. See [DNS Support in Your VPC](#) in the AWS documentation.

If you prefer to use your own Route 53 hosted private zone, you must associate the existing hosted zone with your VPC prior to installing a cluster. You can define your hosted zone using the **platform.aws.hostedZone** and **platform.aws.hostedZoneRole** fields in the **install-config.yaml** file. You can use a private hosted zone from another account by sharing it with the account where you install the cluster. If you use a private hosted zone from another account, you must use the **Passthrough** or **Manual** credentials mode.

If you are working in a disconnected environment, you are unable to reach the public IP addresses for EC2, ELB, and S3 endpoints. Depending on the level to which you want to restrict internet traffic during the installation, the following configuration options are available:

Option 1: Create VPC endpoints

Create a VPC endpoint and attach it to the subnets that the clusters are using. Name the endpoints as follows:

- **ec2.<aws_region>.amazonaws.com.cn**
- **elasticloadbalancing.<aws_region>.amazonaws.com**
- **s3.<aws_region>.amazonaws.com**

With this option, network traffic remains private between your VPC and the required AWS services.

Option 2: Create a proxy without VPC endpoints

As part of the installation process, you can configure an HTTP or HTTPS proxy. With this option, internet traffic goes through the proxy to reach the required AWS services.

Option 3: Create a proxy with VPC endpoints

As part of the installation process, you can configure an HTTP or HTTPS proxy with VPC endpoints. Create a VPC endpoint and attach it to the subnets that the clusters are using. Name the endpoints as follows:

- **ec2.<aws_region>.amazonaws.com.cn**
- **elasticloadbalancing.<aws_region>.amazonaws.com**
- **s3.<aws_region>.amazonaws.com**

When configuring the proxy in the **install-config.yaml** file, add these endpoints to the **noProxy** field. With this option, the proxy prevents the cluster from accessing the internet directly. However, network traffic remains private between your VPC and the required AWS services.

Required VPC components

You must provide a suitable VPC and subnets that allow communication to your machines.

Component	AWS type	Description
VPC	<ul style="list-style-type: none"> • AWS::EC2::VPC • AWS::EC2::VPCEndpoint 	You must provide a public VPC for the cluster to use. The VPC uses an endpoint that references the route tables for each subnet to improve communication with the registry that is hosted in S3.

Component	AWS type	Description												
Public subnets	<ul style="list-style-type: none"> ● AWS::EC2::Subnet ● AWS::EC2::SubnetNetworkACLAssociation 	Your VPC must have public subnets for between 1 and 3 availability zones and associate them with appropriate Ingress rules.												
Internet gateway	<ul style="list-style-type: none"> ● AWS::EC2::InternetGateway ● AWS::EC2::VPCGatewayAttachment ● AWS::EC2::RouteTable ● AWS::EC2::Route ● AWS::EC2::SubnetRouteTableAssociation ● AWS::EC2::NatGateway ● AWS::EC2::EIP 	You must have a public internet gateway, with public routes, attached to the VPC. In the provided templates, each public subnet has a NAT gateway with an EIP address. These NAT gateways allow cluster resources, like private subnet instances, to reach the internet and are not required for some restricted network or proxy scenarios.												
Network access control	<ul style="list-style-type: none"> ● AWS::EC2::NetworkACL ● AWS::EC2::NetworkACLEntry 	<p>You must allow the VPC to access the following ports:</p> <table border="1"> <thead> <tr> <th>Port</th> <th>Reason</th> </tr> </thead> <tbody> <tr> <td>80</td> <td>Inbound HTTP traffic</td> </tr> <tr> <td>443</td> <td>Inbound HTTPS traffic</td> </tr> <tr> <td>22</td> <td>Inbound SSH traffic</td> </tr> <tr> <td>1024 - 65535</td> <td>Inbound ephemeral traffic</td> </tr> <tr> <td>0 - 65535</td> <td>Outbound ephemeral traffic</td> </tr> </tbody> </table>	Port	Reason	80	Inbound HTTP traffic	443	Inbound HTTPS traffic	22	Inbound SSH traffic	1024 - 65535	Inbound ephemeral traffic	0 - 65535	Outbound ephemeral traffic
Port	Reason													
80	Inbound HTTP traffic													
443	Inbound HTTPS traffic													
22	Inbound SSH traffic													
1024 - 65535	Inbound ephemeral traffic													
0 - 65535	Outbound ephemeral traffic													
Private subnets	<ul style="list-style-type: none"> ● AWS::EC2::Subnet ● AWS::EC2::RouteTable ● AWS::EC2::SubnetRouteTableAssociation 	Your VPC can have private subnets. The provided CloudFormation templates can create private subnets for between 1 and 3 availability zones. If you use private subnets, you must provide appropriate routes and tables for them.												

6.3.10.4.2. VPC validation

To ensure that the subnets that you provide are suitable, the installation program confirms the following data:

- All the subnets that you specify exist.
- You provide private subnets.
- The subnet CIDRs belong to the machine CIDR that you specified.
- You provide subnets for each availability zone. Each availability zone contains no more than one public and one private subnet. If you use a private cluster, provide only a private subnet for each availability zone. Otherwise, provide exactly one public and private subnet for each availability zone.
- You provide a public subnet for each private subnet availability zone. Machines are not provisioned in availability zones that you do not provide private subnets for.

If you destroy a cluster that uses an existing VPC, the VPC is not deleted. When you remove the OpenShift Container Platform cluster from a VPC, the **kubernetes.io/cluster/.*: shared** tag is removed from the subnets that it used.

6.3.10.4.3. Division of permissions

Starting with OpenShift Container Platform 4.3, you do not need all of the permissions that are required for an installation program-provisioned infrastructure cluster to deploy a cluster. This change mimics the division of permissions that you might have at your company: some individuals can create different resource in your clouds than others. For example, you might be able to create application-specific items, like instances, buckets, and load balancers, but not networking-related components such as VPCs, subnets, or ingress rules.

The AWS credentials that you use when you create your cluster do not need the networking permissions that are required to make VPCs and core networking components within the VPC, such as subnets, routing tables, internet gateways, NAT, and VPN. You still need permission to make the application resources that the machines within the cluster require, such as ELBs, security groups, S3 buckets, and nodes.

6.3.10.4.4. Isolation between clusters

If you deploy OpenShift Container Platform to an existing network, the isolation of cluster services is reduced in the following ways:

- You can install multiple OpenShift Container Platform clusters in the same VPC.
- ICMP ingress is allowed from the entire network.
- TCP 22 ingress (SSH) is allowed to the entire network.
- Control plane TCP 6443 ingress (Kubernetes API) is allowed to the entire network.
- Control plane TCP 22623 ingress (MCS) is allowed to the entire network.

6.3.10.4.5. Optional: AWS security groups

By default, the installation program creates and attaches security groups to control plane and compute machines. The rules associated with the default security groups cannot be modified.

However, you can apply additional existing AWS security groups, which are associated with your existing VPC, to control plane and compute machines. Applying custom security groups can help you meet the security needs of your organization, in such cases where you need to control the incoming or outgoing traffic of these machines.

As part of the installation process, you apply custom security groups by modifying the **install-config.yaml** file before deploying the cluster.

For more information, see "Applying existing AWS security groups to the cluster".

6.3.10.5. Uploading a custom RHCOS AMI in AWS

If you are deploying to a custom Amazon Web Services (AWS) region, you must upload a custom Red Hat Enterprise Linux CoreOS (RHCOS) Amazon Machine Image (AMI) that belongs to that region.

Prerequisites

- You configured an AWS account.
- You created an Amazon S3 bucket with the required IAM [service role](#).
- You uploaded your RHCOS VMDK file to Amazon S3. The RHCOS VMDK file must be the highest version that is less than or equal to the OpenShift Container Platform version you are installing.
- You downloaded the AWS CLI and installed it on your computer. See [Install the AWS CLI Using the Bundled Installer](#).

Procedure

1. Export your AWS profile as an environment variable:

```
$ export AWS_PROFILE=<aws_profile> 1
```

- 1 The AWS profile name that holds your AWS credentials, like **beijingadmin**.

2. Export the region to associate with your custom AMI as an environment variable:

```
$ export AWS_DEFAULT_REGION=<aws_region> 1
```

- 1 The AWS region, like **cn-north-1**.

3. Export the version of RHCOS you uploaded to Amazon S3 as an environment variable:

```
$ export RHCOS_VERSION=<version> 1
```

- 1 The RHCOS VMDK version, like **4.16.0**.

4. Export the Amazon S3 bucket name as an environment variable:

```
$ export VMIMPORT_BUCKET_NAME=<s3_bucket_name>
```

5. Create the **containers.json** file and define your RHCOS VMDK file:

```
$ cat <<EOF > containers.json
{
  "Description": "rhcos- $\{RHCOS\_VERSION\}$ -x86_64-aws.x86_64",
  "Format": "vmdk",
  "UserBucket": {
    "S3Bucket": " $\{VMIMPORT\_BUCKET\_NAME\}$ ",
    "S3Key": "rhcos- $\{RHCOS\_VERSION\}$ -x86_64-aws.x86_64.vmdk"
  }
}
EOF
```

6. Import the RHCOS disk as an Amazon EBS snapshot:

```
$ aws ec2 import-snapshot --region  $\{AWS\_DEFAULT\_REGION\}$  \
  --description "<description>" \ ❶
  --disk-container "file://<file_path>/containers.json" ❷
```

- ❶ The description of your RHCOS disk being imported, like **rhcos- $\{RHCOS_VERSION\}$ -x86_64-aws.x86_64**.
- ❷ The file path to the JSON file describing your RHCOS disk. The JSON file should contain your Amazon S3 bucket name and key.

7. Check the status of the image import:

```
$ watch -n 5 aws ec2 describe-import-snapshot-tasks --region  $\{AWS\_DEFAULT\_REGION\}$ 
```

Example output

```
{
  "ImportSnapshotTasks": [
    {
      "Description": "rhcos-4.7.0-x86_64-aws.x86_64",
      "ImportTaskId": "import-snap-fh6i8uil",
      "SnapshotTaskDetail": {
        "Description": "rhcos-4.7.0-x86_64-aws.x86_64",
        "DiskImageSize": 819056640.0,
        "Format": "VMDK",
        "SnapshotId": "snap-06331325870076318",
        "Status": "completed",
        "UserBucket": {
          "S3Bucket": "external-images",
          "S3Key": "rhcos-4.7.0-x86_64-aws.x86_64.vmdk"
        }
      }
    }
  ]
}
```

Copy the **SnapshotId** to register the image.

8. Create a custom RHCOS AMI from the RHCOS snapshot:

```
$ aws ec2 register-image \
  --region ${AWS_DEFAULT_REGION} \
  --architecture x86_64 \ 1
  --description "rhcos-${RHCOS_VERSION}-x86_64-aws.x86_64" \ 2
  --ena-support \
  --name "rhcos-${RHCOS_VERSION}-x86_64-aws.x86_64" \ 3
  --virtualization-type hvm \
  --root-device-name '/dev/xvda' \
  --block-device-mappings 'DeviceName=/dev/xvda,Ebs=
  {DeleteOnTermination=true,SnapshotId=<snapshot_ID>' 4
```

- 1** The RHCOS VMDK architecture type, like **x86_64**, **aarch64**, **s390x**, or **ppc64le**.
- 2** The **Description** from the imported snapshot.
- 3** The name of the RHCOS AMI.
- 4** The **SnapshotID** from the imported snapshot.

To learn more about these APIs, see the AWS documentation for [importing snapshots](#) and [creating EBS-backed AMIs](#).

6.3.10.6. Manually creating the installation configuration file

Installing the cluster requires that you manually create the installation configuration file.


Prerequisites

- You have uploaded a custom RHCOS AMI.
- You have an SSH public key on your local machine to provide to the installation program. The key will be used for SSH authentication onto your cluster nodes for debugging and disaster recovery.
- You have obtained the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create an installation directory to store your required installation assets in:


```
$ mkdir <installation_directory>
```



IMPORTANT


You must create a directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

2. Customize the sample **install-config.yaml** file template that is provided and save it in the **<installation_directory>**.

**NOTE**

You must name this configuration file **install-config.yaml**.

3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.

**IMPORTANT**


The **install-config.yaml** file is consumed during the next step of the installation process. You must back it up now.

Additional resources

- [Installation configuration parameters for AWS](#)

6.3.10.6.1. Sample customized install-config.yaml file for AWS

You can customize the installation configuration file (**install-config.yaml**) to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.

**IMPORTANT**

This sample YAML file is provided for reference only. Use it as a resource to enter parameter values into the installation configuration file that you created manually.

```

apiVersion: v1
baseDomain: example.com 1
credentialsMode: Mint 2
controlPlane: 3 4
  hyperthreading: Enabled 5
  name: master
  platform:
    aws:
      zones:
      - cn-north-1a
      - cn-north-1b
      rootVolume:
        iops: 4000
        size: 500
        type: io1 6
      metadataService:
        authentication: Optional 7
        type: m6i.xlarge
      replicas: 3
  compute: 8
  - hyperthreading: Enabled 9
    name: worker
    platform:
      aws:

```

```

rootVolume:
  iops: 2000
  size: 500
  type: io1 10
metadataService:
  authentication: Optional 11
  type: c5.4xlarge
  zones:
    - cn-north-1a
replicas: 3
metadata:
  name: test-cluster 12
networking:
  clusterNetwork:
    - cidr: 10.128.0.0/14
    hostPrefix: 23
  machineNetwork:
    - cidr: 10.0.0.0/16
  networkType: OVNKubernetes 13
  serviceNetwork:
    - 172.30.0.0/16
platform:
  aws:
    region: cn-north-1 14
    propagateUserTags: true 15
    userTags:
      adminContact: jdoe
      costCenter: 7536
    subnets: 16
    - subnet-1
    - subnet-2
    - subnet-3
    amiID: ami-96c6f8f7 17 18
    serviceEndpoints: 19
    - name: ec2
      url: https://vpce-id.ec2.cn-north-1.vpce.amazonaws.com.cn
    hostedZone: Z3URY6TWQ91KVV 20
  fips: false 21
  sshKey: ssh-ed25519 AAAA... 22
  publish: Internal 23
  pullSecret: '{"auths": ...}' 24

```

1 **12** **14** **17** **24** Required.

2 Optional: Add this parameter to force the Cloud Credential Operator (CCO) to use the specified mode. By default, the CCO uses the root credentials in the **kube-system** namespace to dynamically try to determine the capabilities of the credentials. For details about CCO modes, see the "About the Cloud Credential Operator" section in the *Authentication and authorization* guide.

3 **8** **15** If you do not provide these parameters and values, the installation program provides the default value.

4

The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section

- 5 9** Whether to enable or disable simultaneous multithreading, or **hyperthreading**. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores. You can disable it by setting the parameter value to **Disabled**. If you disable simultaneous multithreading in some cluster machines, you must disable it in all cluster machines.



IMPORTANT

If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance. Use larger instance types, such as **m4.2xlarge** or **m5.2xlarge**, for your machines if you disable simultaneous multithreading.

- 6 10** To configure faster storage for etcd, especially for larger clusters, set the storage type as **io1** and set **iops** to **2000**.

- 7 11** Whether to require the [Amazon EC2 Instance Metadata Service v2](#) (IMDSv2). To require IMDSv2, set the parameter value to **Required**. To allow the use of both IMDSv1 and IMDSv2, set the parameter value to **Optional**. If no value is specified, both IMDSv1 and IMDSv2 are allowed.



NOTE

The IMDS configuration for control plane machines that is set during cluster installation can only be changed by using the AWS CLI. The IMDS configuration for compute machines can be changed by using compute machine sets.

- 13** The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 16** If you provide your own VPC, specify subnets for each availability zone that your cluster uses.
- 18** The ID of the AMI used to boot machines for the cluster. If set, the AMI must belong to the same region as the cluster.
- 19** The AWS service endpoints. Custom endpoints are required when installing to an unknown AWS region. The endpoint URL must use the **https** protocol and the host must trust the certificate.
- 20** The ID of your existing Route 53 private hosted zone. Providing an existing hosted zone requires that you supply your own VPC and the hosted zone is already associated with the VPC prior to installing your cluster. If undefined, the installation program creates a new hosted zone.
- 21** Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.



IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 22** You can optionally provide the **sshKey** value that you use to access the machines in your cluster.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- 23** How to publish the user-facing endpoints of your cluster. Set **publish** to **Internal** to deploy a private cluster, which cannot be accessed from the internet. The default value is **External**.

6.3.10.6.2. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 6.18. Minimum resource requirements

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Bootstrap	RHCOS	4	16 GB	100 GB	300
Control plane	RHCOS	4	16 GB	100 GB	300
Compute	RHCOS, RHEL 8.6 and later [3]	2	8 GB	100 GB	300

1. One vCPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio: (threads per core × cores) × sockets = vCPUs.
2. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes which require a 10 ms p99 fsync duration. Note that on many cloud platforms, storage size and IOPS scale together, so you might need to over-allocate storage volume to obtain sufficient performance.
3. As with all user-provisioned installations, if you choose to use RHEL compute machines in your

cluster, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Use of RHEL 7 compute machines is deprecated and has been removed in OpenShift Container Platform 4.10 and later.



NOTE

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

6.3.10.6.3. Tested instance types for AWS

The following Amazon Web Services (AWS) instance types have been tested with OpenShift Container Platform.



NOTE

Use the machine types included in the following charts for your AWS instances. If you use an instance type that is not listed in the chart, ensure that the instance size you use matches the minimum resource requirements that are listed in the section named "Minimum resource requirements for cluster installation".


Example 6.44. Machine types based on 64-bit x86 architecture

- **c4.***
- **c5.***
- **c5a.***
- **i3.***
- **m4.***
- **m5.***
- **m5a.***

- **m6a.***
- **m6i.***
- **r4.***
- **r5.***
- **r5a.***
- **r6i.***
- **t3.***
- **t3a.***

6.3.10.6.4. Tested instance types for AWS on 64-bit ARM infrastructures

The following Amazon Web Services (AWS) 64-bit ARM instance types have been tested with OpenShift Container Platform.



NOTE

Use the machine types included in the following charts for your AWS ARM instances. If you use an instance type that is not listed in the chart, ensure that the instance size you use matches the minimum resource requirements that are listed in "Minimum resource requirements for cluster installation".

Example 6.45. Machine types based on 64-bit ARM architecture


- **c6g.***
- **m6g.***

6.3.10.6.5. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: ec2.<aws_region>.amazonaws.com,elasticloadbalancing.
<aws_region>.amazonaws.com,s3.<aws_region>.amazonaws.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations. If you have added the Amazon **EC2**, **Elastic Load Balancing**, and **S3** VPC endpoints to your VPC, you must add these endpoints to the **noProxy** field.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.

**NOTE**

The installation program does not support the proxy **readinessEndpoints** field.


**NOTE**

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.

**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

6.3.10.6.6. Applying existing AWS security groups to the cluster

Applying existing AWS security groups to your control plane and compute machines can help you meet the security needs of your organization, in such cases where you need to control the incoming or outgoing traffic of these machines.

Prerequisites

- You have created the security groups in AWS. For more information, see the AWS documentation about working with [security groups](#).
- The security groups must be associated with the existing VPC that you are deploying the cluster to. The security groups cannot be associated with another VPC.
- You have an existing **install-config.yaml** file.

Procedure

1. In the **install-config.yaml** file, edit the **compute.platform.aws.additionalSecurityGroupIDs** parameter to specify one or more custom security groups for your compute machines.
2. Edit the **controlPlane.platform.aws.additionalSecurityGroupIDs** parameter to specify one or more custom security groups for your control plane machines.
3. Save the file and reference it when deploying the cluster.

Sample **install-config.yaml** file that specifies custom security groups

```
# ...
compute:
- hyperthreading: Enabled
  name: worker
```

```

platform:
  aws:
    additionalSecurityGroupIDs:
      - sg-1 1
      - sg-2
  replicas: 3
controlPlane:
  hyperthreading: Enabled
  name: master
platform:
  aws:
    additionalSecurityGroupIDs:
      - sg-3
      - sg-4
  replicas: 3
platform:
  aws:
    region: us-east-1
  subnets: 2
    - subnet-1
    - subnet-2
    - subnet-3

```

- 1** Specify the name of the security group as it appears in the Amazon EC2 console, including the **sg** prefix.
- 2** Specify subnets for each availability zone that your cluster uses.

6.3.10.7. Alternatives to storing administrator-level secrets in the kube-system project

By default, administrator secrets are stored in the **kube-system** project. If you configured the **credentialsMode** parameter in the **install-config.yaml** file to **Manual**, you must use one of the following alternatives:

- To manage long-term cloud credentials manually, follow the procedure in [Manually creating long-term credentials](#).
- To implement short-term credentials that are managed outside the cluster for individual components, follow the procedures in [Configuring an AWS cluster to use short-term credentials](#).

6.3.10.7.1. Manually creating long-term credentials

The Cloud Credential Operator (CCO) can be put into manual mode prior to installation in environments where the cloud identity and access management (IAM) APIs are not reachable, or the administrator prefers not to store an administrator-level credential secret in the cluster **kube-system** namespace.

Procedure

1. If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```

apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...

```

- If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

- Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

- Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```

$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included \ 1
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \ 2
  --to=<path_to_directory_for_credentials_requests> \ 3

```

- The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- Specify the location of the **install-config.yaml** file.
- Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

This command creates a YAML file for each **CredentialsRequest** object.

Sample CredentialsRequest object

```

apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: AWSProviderSpec
    statementEntries:
      - effect: Allow
        action:
          - iam:GetUser

```

```

- iam:GetUserPolicy
- iam:ListAccessKeys
resource: "*"
...

```

5. Create YAML files for secrets in the **openshift-install** manifests directory that you generated previously. The secrets must be stored using the namespace and secret name defined in the **spec.secretRef** for each **CredentialsRequest** object.

Sample CredentialsRequest object with secrets

```

apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: AWSProviderSpec
    statementEntries:
      - effect: Allow
        action:
          - s3:CreateBucket
          - s3>DeleteBucket
        resource: "*"
      ...
  secretRef:
    name: <component_secret>
    namespace: <component_namespace>
  ...

```

Sample Secret object

```

apiVersion: v1
kind: Secret
metadata:
  name: <component_secret>
  namespace: <component_namespace>
data:
  aws_access_key_id: <base64_encoded_aws_access_key_id>
  aws_secret_access_key: <base64_encoded_aws_secret_access_key>

```



IMPORTANT

Before upgrading a cluster that uses manually maintained credentials, you must ensure that the CCO is in an upgradeable state.

6.3.10.7.2. Configuring an AWS cluster to use short-term credentials

To install a cluster that is configured to use the AWS Security Token Service (STS), you must configure the CCO utility and create the required AWS resources for your cluster.

6.3.10.7.2.1. Configuring the Cloud Credential Operator utility

To create and manage cloud credentials from outside of the cluster when the Cloud Credential Operator (CCO) is operating in manual mode, extract and prepare the CCO utility (**ccoctl**) binary.



NOTE

The **ccoctl** utility is a Linux binary that must run in a Linux environment.

Prerequisites

- You have access to an OpenShift Container Platform account with cluster administrator access.
- You have installed the OpenShift CLI (**oc**).
- You have created an AWS account for the **ccoctl** utility to use with the following permissions:

Example 6.46. Required AWS permissions

Required iam permissions

- **iam:CreateOpenIDConnectProvider**
- **iam:CreateRole**
- **iam>DeleteOpenIDConnectProvider**
- **iam>DeleteRole**
- **iam>DeleteRolePolicy**
- **iam:GetOpenIDConnectProvider**
- **iam:GetRole**
- **iam:GetUser**
- **iam:ListOpenIDConnectProviders**
- **iam:ListRolePolicies**
- **iam:ListRoles**
- **iam:PutRolePolicy**
- **iam:TagOpenIDConnectProvider**
- **iam:TagRole**

Required s3 permissions

- **s3:CreateBucket**
- **s3>DeleteBucket**
- **s3>DeleteObject**
- **s3:GetBucketAcl**

- **s3:GetBucketTagging**
- **s3:GetObject**
- **s3:GetObjectAcl**
- **s3:GetObjectTagging**
- **s3:ListBucket**
- **s3:PutBucketAcl**
- **s3:PutBucketPolicy**
- **s3:PutBucketPublicAccessBlock**
- **s3:PutBucketTagging**
- **s3:PutObject**
- **s3:PutObjectAcl**
- **s3:PutObjectTagging**

Required cloudfront permissions

- **cloudfront:ListCloudFrontOriginAccessIdentities**
- **cloudfront:ListDistributions**
- **cloudfront:ListTagsForResource**

If you plan to store the OIDC configuration in a private S3 bucket that is accessed by the IAM identity provider through a public CloudFront distribution URL, the AWS account that runs the **ccoctl** utility requires the following additional permissions:

Example 6.47. Additional permissions for a private S3 bucket with CloudFront

- **cloudfront:CreateCloudFrontOriginAccessIdentity**
- **cloudfront:CreateDistribution**
- **cloudfront>DeleteCloudFrontOriginAccessIdentity**
- **cloudfront>DeleteDistribution**
- **cloudfront:GetCloudFrontOriginAccessIdentity**
- **cloudfront:GetCloudFrontOriginAccessIdentityConfig**
- **cloudfront:GetDistribution**
- **cloudfront:TagResource**
- **cloudfront:UpdateDistribution**

**NOTE**

These additional permissions support the use of the **--create-private-s3-bucket** option when processing credentials requests with the **ccoctl aws create-all** command.

Procedure

1. Set a variable for the OpenShift Container Platform release image by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

2. Obtain the CCO container image from the OpenShift Container Platform release image by running the following command:

```
$ CCO_IMAGE=$(oc adm release info --image-for='cloud-credential-operator'
$RELEASE_IMAGE -a ~/.pull-secret)
```

**NOTE**

Ensure that the architecture of the **\$RELEASE_IMAGE** matches the architecture of the environment in which you will use the **ccoctl** tool.

3. Extract the **ccoctl** binary from the CCO container image within the OpenShift Container Platform release image by running the following command:

```
$ oc image extract $CCO_IMAGE \
--file="/usr/bin/ccoctl.<rhel_version>" \
-a ~/.pull-secret
```

- 1 For **<rhel_version>**, specify the value that corresponds to the version of Red Hat Enterprise Linux (RHEL) that the host uses. If no value is specified, **ccoctl.rhel8** is used by default. The following values are valid:

- **rhel8**: Specify this value for hosts that use RHEL 8.
- **rhel9**: Specify this value for hosts that use RHEL 9.

4. Change the permissions to make **ccoctl** executable by running the following command:

```
$ chmod 775 ccoctl.<rhel_version>
```

Verification

- To verify that **ccoctl** is ready to use, display the help file by running the following command:

```
$ ccoctl --help
```

Output of ccoctl --help

OpenShift credentials provisioning tool

Usage:

`ccoctl [command]`

Available Commands:

`aws` Manage credentials objects for AWS cloud
`azure` Manage credentials objects for Azure
`gcp` Manage credentials objects for Google cloud
`help` Help about any command
`ibmcloud` Manage credentials objects for IBM Cloud
`nutanix` Manage credentials objects for Nutanix

Flags:

`-h, --help` help for `ccoctl`

Use "`ccoctl [command] --help`" for more information about a command.

6.3.10.7.2.2. Creating AWS resources with the Cloud Credential Operator utility


You have the following options when creating AWS resources:

- You can use the **`ccoctl aws create-all`** command to create the AWS resources automatically. This is the quickest way to create the resources. See [Creating AWS resources with a single command](#).
- If you need to review the JSON files that the **`ccoctl`** tool creates before modifying AWS resources, or if the process the **`ccoctl`** tool uses to create AWS resources automatically does not meet the requirements of your organization, you can create the AWS resources individually. See [Creating AWS resources individually](#).

6.3.10.7.2.2.1. Creating AWS resources with a single command

If the process the **`ccoctl`** tool uses to create AWS resources automatically meets the requirements of your organization, you can use the **`ccoctl aws create-all`** command to automate the creation of AWS resources.

Otherwise, you can create the AWS resources individually. For more information, see "Creating AWS resources individually".



NOTE

By default, **`ccoctl`** creates objects in the directory in which the commands are run. To create the objects in a different directory, use the **`--output-dir`** flag. This procedure uses **`<path_to_ccoctl_output_dir>`** to refer to this directory.

Prerequisites

You must have:

- Extracted and prepared the **`ccoctl`** binary.

Procedure

1. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

2. Extract the list of **CredentialsRequest** objects from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included \1
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \2
  --to=<path_to_directory_for_credentials_requests> \3
```

- 1 The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2 Specify the location of the **install-config.yaml** file.
- 3 Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.



NOTE

This command might take a few moments to run.

3. Use the **ccoctl** tool to process all **CredentialsRequest** objects by running the following command:

```
$ ccoctl aws create-all \
  --name=<name> \1
  --region=<aws_region> \2
  --credentials-requests-dir=<path_to_credentials_requests_directory> \3
  --output-dir=<path_to_ccoctl_output_dir> \4
  --create-private-s3-bucket \5
```

- 1 Specify the name used to tag any cloud resources that are created for tracking.
- 2 Specify the AWS region in which cloud resources will be created.
- 3 Specify the directory containing the files for the component **CredentialsRequest** objects.
- 4 Optional: Specify the directory in which you want the **ccoctl** utility to create objects. By default, the utility creates objects in the directory in which the commands are run.
- 5 Optional: By default, the **ccoctl** utility stores the OpenID Connect (OIDC) configuration files in a public S3 bucket and uses the S3 URL as the public OIDC endpoint. To store the OIDC configuration in a private S3 bucket that is accessed by the IAM identity provider through a public CloudFront distribution URL instead, use the **--create-private-s3-bucket** parameter.

**NOTE**

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

Verification

- To verify that the OpenShift Container Platform secrets are created, list the files in the **<path_to_ccoctl_output_dir>/manifests** directory:

```
$ ls <path_to_ccoctl_output_dir>/manifests
```

Example output

```
cluster-authentication-02-config.yaml
openshift-cloud-credential-operator-cloud-credential-operator-iam-ro-creds-credentials.yaml
openshift-cloud-network-config-controller-cloud-credentials-credentials.yaml
openshift-cluster-api-capamanager-bootstrap-credentials-credentials.yaml
openshift-cluster-csi-drivers-ebs-cloud-credentials-credentials.yaml
openshift-image-registry-installer-cloud-credentials-credentials.yaml
openshift-ingress-operator-cloud-credentials-credentials.yaml
openshift-machine-api-aws-cloud-credentials-credentials.yaml
```

You can verify that the IAM roles are created by querying AWS. For more information, refer to AWS documentation on listing IAM roles.

6.3.10.7.2.2.2. Creating AWS resources individually

You can use the **ccoctl** tool to create AWS resources individually. This option might be useful for an organization that shares the responsibility for creating these resources among different users or departments.

Otherwise, you can use the **ccoctl aws create-all** command to create the AWS resources automatically. For more information, see "Creating AWS resources with a single command".

**NOTE**

By default, **ccoctl** creates objects in the directory in which the commands are run. To create the objects in a different directory, use the **--output-dir** flag. This procedure uses **<path_to_ccoctl_output_dir>** to refer to this directory.

Some **ccoctl** commands make AWS API calls to create or modify AWS resources. You can use the **--dry-run** flag to avoid making API calls. Using this flag creates JSON files on the local file system instead. You can review and modify the JSON files and then apply them with the AWS CLI tool using the **--cli-input-json** parameters.

Prerequisites

- Extract and prepare the **ccoctl** binary.

Procedure

1. Generate the public and private RSA key files that are used to set up the OpenID Connect provider for the cluster by running the following command:

```
$ ccoctl aws create-key-pair
```

Example output

```
2021/04/13 11:01:02 Generating RSA keypair
2021/04/13 11:01:03 Writing private key to /<path_to_ccoctl_output_dir>/serviceaccount-
signer.private
2021/04/13 11:01:03 Writing public key to /<path_to_ccoctl_output_dir>/serviceaccount-
signer.public
2021/04/13 11:01:03 Copying signing key for use by installer
```

where **serviceaccount-signer.private** and **serviceaccount-signer.public** are the generated key files.

This command also creates a private key that the cluster requires during installation in **/<path_to_ccoctl_output_dir>/tls/bound-service-account-signing-key.key**.

2. Create an OpenID Connect identity provider and S3 bucket on AWS by running the following command:

```
$ ccoctl aws create-identity-provider \
  --name=<name> \ 1
  --region=<aws_region> \ 2
  --public-key-file=<path_to_ccoctl_output_dir>/serviceaccount-signer.public 3
```

1 **<name>** is the name used to tag any cloud resources that are created for tracking.

2 **<aws-region>** is the AWS region in which cloud resources will be created.

3 **<path_to_ccoctl_output_dir>** is the path to the public key file that the **ccoctl aws create-key-pair** command generated.

Example output

```
2021/04/13 11:16:09 Bucket <name>-oidc created
2021/04/13 11:16:10 OpenID Connect discovery document in the S3 bucket <name>-oidc at
.well-known/openid-configuration updated
2021/04/13 11:16:10 Reading public key
2021/04/13 11:16:10 JSON web key set (JWKS) in the S3 bucket <name>-oidc at keys.json
updated
2021/04/13 11:16:18 Identity Provider created with ARN: arn:aws:iam::
<aws_account_id>:oidc-provider/<name>-oidc.s3.<aws_region>.amazonaws.com
```

where **openid-configuration** is a discovery document and **keys.json** is a JSON web key set file.

This command also creates a YAML configuration file in **/<path_to_ccoctl_output_dir>/manifests/cluster-authentication-02-config.yaml**. This file sets the issuer URL field for the service account tokens that the cluster generates, so that the AWS IAM identity provider trusts the tokens.

3. Create IAM roles for each component in the cluster:

- a. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

- b. Extract the list of **CredentialsRequest** objects from the OpenShift Container Platform release image:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included 1 \
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \
2 \
  --to=<path_to_directory_for_credentials_requests> 3
```

- 1** The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2** Specify the location of the **install-config.yaml** file.
- 3** Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

- c. Use the **ccoctl** tool to process all **CredentialsRequest** objects by running the following command:

```
$ ccoctl aws create-iam-roles \
  --name=<name> \
  --region=<aws_region> \
  --credentials-requests-dir=<path_to_credentials_requests_directory> \
  --identity-provider-arn=arn:aws:iam::<aws_account_id>:oidc-provider/<name>-oidc.s3.
  <aws_region>.amazonaws.com
```

**NOTE**

For AWS environments that use alternative IAM API endpoints, such as GovCloud, you must also specify your region with the **--region** parameter.

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

For each **CredentialsRequest** object, **ccoctl** creates an IAM role with a trust policy that is tied to the specified OIDC identity provider, and a permissions policy as defined in each **CredentialsRequest** object from the OpenShift Container Platform release image.

Verification

- To verify that the OpenShift Container Platform secrets are created, list the files in the **<path_to_ccoctl_output_dir>/manifests** directory:

```
$ ls <path_to_ccoctl_output_dir>/manifests
```

Example output

```
cluster-authentication-02-config.yaml
openshift-cloud-credential-operator-cloud-credential-operator-iam-ro-creds-credentials.yaml
openshift-cloud-network-config-controller-cloud-credentials-credentials.yaml
openshift-cluster-api-capamanager-bootstrap-credentials-credentials.yaml
openshift-cluster-csi-drivers-ebs-cloud-credentials-credentials.yaml
openshift-image-registry-installer-cloud-credentials-credentials.yaml
openshift-ingress-operator-cloud-credentials-credentials.yaml
openshift-machine-api-aws-cloud-credentials-credentials.yaml
```

You can verify that the IAM roles are created by querying AWS. For more information, refer to AWS documentation on listing IAM roles.

6.3.10.7.2.3. Incorporating the Cloud Credential Operator utility manifests

To implement short-term security credentials managed outside the cluster for individual components, you must move the manifest files that the Cloud Credential Operator utility (**ccoctl**) created to the correct directories for the installation program.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have configured the Cloud Credential Operator utility (**ccoctl**).
- You have created the cloud provider resources that are required for your cluster with the **ccoctl** utility.

Procedure

1. If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```
apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...
```

2. If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

3. Copy the manifests that the **ccoctl** utility generated to the **manifests** directory that the installation program created by running the following command:

```
$ cp /<path_to_ccoctl_output_dir>/manifests/* ./manifests/
```

4. Copy the private key that the **ccoctl** utility generated in the **tls** directory to the installation directory by running the following command:

```
$ cp -a /<path_to_ccoctl_output_dir>/tls .
```

6.3.10.8. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.

Procedure

1. Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1  
--log-level=info 2
```

- 1** For **<installation_directory>**, specify the location of your customized **./install-config.yaml** file.
- 2** To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

2. Optional: Remove or disable the **AdministratorAccess** policy from the IAM account that you used to install the cluster.



NOTE

The elevated permissions provided by the **AdministratorAccess** policy are required only during installation.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to **<installation_directory>/openshift_install.log**.



IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

6.3.10.9. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

6.3.10.10. Logging in to the cluster by using the web console

The **kubeadmin** user exists by default after an OpenShift Container Platform installation. You can log in to your cluster as the **kubeadmin** user by using the OpenShift Container Platform web console.

Prerequisites

- You have access to the installation host.
- You completed a cluster installation and all cluster Operators are available.

Procedure

1. Obtain the password for the **kubeadmin** user from the **kubeadmin-password** file on the installation host:

```
$ cat <installation_directory>/auth/kubeadmin-password
```



NOTE

Alternatively, you can obtain the **kubeadmin** password from the **<installation_directory>/openshift_install.log** log file on the installation host.

2. List the OpenShift Container Platform web console route:

```
$ oc get routes -n openshift-console | grep 'console-openshift'
```



NOTE

Alternatively, you can obtain the OpenShift Container Platform route from the **<installation_directory>/openshift_install.log** log file on the installation host.

Example output

```
console    console-openshift-console.apps.<cluster_name>.<base_domain>    console
https reencrypt/Redirect None
```


3. Navigate to the route detailed in the output of the preceding command in a web browser and log in as the **kubeadmin** user.

Additional resources

- See [Accessing the web console](#) for more details about accessing and understanding the OpenShift Container Platform web console.

6.3.10.11. Next steps

- [Validating an installation](#).
- [Customize your cluster](#).
- If necessary, you can [opt out of remote health reporting](#).
- If necessary, you can [remove cloud provider credentials](#).

6.3.11. Installing a cluster with compute nodes on AWS Local Zones

You can quickly install an OpenShift Container Platform cluster on Amazon Web Services (AWS) Local Zones by setting the zone names in the edge compute pool of the **install-config.yaml** file, or install a cluster in an existing Amazon Virtual Private Cloud (VPC) with Local Zone subnets.

AWS Local Zones is an infrastructure that place Cloud Resources close to metropolitan regions. For more information, see the [AWS Local Zones Documentation](#).

6.3.11.1. Infrastructure prerequisites

- You reviewed details about [OpenShift Container Platform installation and update](#) processes.
- You are familiar with [Selecting a cluster installation method and preparing it for users](#).
- You [configured an AWS account](#) to host the cluster.



WARNING

If you have an AWS profile stored on your computer, it must not use a temporary session token that you generated while using a multi-factor authentication device. The cluster continues to use your current AWS credentials to create AWS resources for the entire life of the cluster, so you must use key-based, long-term credentials. To generate appropriate keys, see [Managing Access Keys for IAM Users](#) in the AWS documentation. You can supply the keys when you run the installation program.

- You downloaded the AWS CLI and installed it on your computer. See [Install the AWS CLI Using the Bundled Installer \(Linux, macOS, or UNIX\)](#) in the AWS documentation.
- If you use a firewall, you [configured it to allow the sites](#) that your cluster must access.

- You noted the region and supported [AWS Local Zones locations](#) to create the network resources in.
- You read the [AWS Local Zones features](#) in the AWS documentation.
- You added permissions for creating network resources that support AWS Local Zones to the Identity and Access Management (IAM) user or role. The following example enables a zone group that can provide a user or role access for creating network resources that support AWS Local Zones.

Example of an additional IAM policy with the `ec2:ModifyAvailabilityZoneGroup` permission attached to an IAM user or role.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "ec2:ModifyAvailabilityZoneGroup"
      ],
      "Effect": "Allow",
      "Resource": "*"
    }
  ]
}
```

6.3.11.2. About AWS Local Zones and edge compute pool

Read the following sections to understand infrastructure behaviors and cluster limitations in an AWS Local Zones environment.

6.3.11.2.1. Cluster limitations in AWS Local Zones

Some limitations exist when you try to deploy a cluster with a default installation configuration in an Amazon Web Services (AWS) Local Zone.



IMPORTANT

The following list details limitations when deploying a cluster in a pre-configured AWS zone:

- The maximum transmission unit (MTU) between an Amazon EC2 instance in a zone and an Amazon EC2 instance in the Region is **1300**. This causes the cluster-wide network MTU to change according to the network plugin that is used with the deployment.
- Network resources such as Network Load Balancer (NLB), Classic Load Balancer, and Network Address Translation (NAT) Gateways are not globally supported.
- For an OpenShift Container Platform cluster on AWS, the AWS Elastic Block Storage (EBS) **gp3** type volume is the default for node volumes and the default for the storage class. This volume type is not globally available on zone locations. By default, the nodes running in zones are deployed with the **gp2** EBS volume. The **gp2-csi StorageClass** parameter must be set when creating workloads on zone nodes.

If you want the installation program to automatically create Local Zone subnets for your OpenShift Container Platform cluster, specific configuration limitations apply with this method.



IMPORTANT

The following configuration limitation applies when you set the installation program to automatically create subnets for your OpenShift Container Platform cluster:

- When the installation program creates private subnets in AWS Local Zones, the program associates each subnet with the route table of its parent zone. This operation ensures that each private subnet can route egress traffic to the internet by way of NAT Gateways in an AWS Region.
- If the parent-zone route table does not exist during cluster installation, the installation program associates any private subnet with the first available private route table in the Amazon Virtual Private Cloud (VPC). This approach is valid only for AWS Local Zones subnets in an OpenShift Container Platform cluster.

6.3.11.2.2. About edge compute pools

Edge compute nodes are tainted compute nodes that run in AWS Local Zones locations.

When deploying a cluster that uses Local Zones, consider the following points:

- Amazon EC2 instances in the Local Zones are more expensive than Amazon EC2 instances in the Availability Zones.
- The latency is lower between the applications running in AWS Local Zones and the end user. A latency impact exists for some workloads if, for example, ingress traffic is mixed between Local Zones and Availability Zones.



IMPORTANT

Generally, the maximum transmission unit (MTU) between an Amazon EC2 instance in a Local Zones and an Amazon EC2 instance in the Region is 1300. The cluster network MTU must be always less than the EC2 MTU to account for the overhead. The specific overhead is determined by the network plugin. For example: OVN-Kubernetes has an overhead of **100 bytes**.

The network plugin can provide additional features, such as IPsec, that also affect the MTU sizing.

For more information, see [How Local Zones work](#) in the AWS documentation.

OpenShift Container Platform 4.12 introduced a new compute pool, *edge*, that is designed for use in remote zones. The edge compute pool configuration is common between AWS Local Zones locations. Because of the type and size limitations of resources like EC2 and EBS on Local Zones resources, the default instance type can vary from the traditional compute pool.

The default Elastic Block Store (EBS) for Local Zones locations is **gp2**, which differs from the non-edge compute pool. The instance type used for each Local Zones on an edge compute pool also might differ from other compute pools, depending on the instance offerings on the zone.

The edge compute pool creates new labels that developers can use to deploy applications onto AWS Local Zones nodes. The new labels are:

- **node-role.kubernetes.io/edge=""**
- **machine.openshift.io/zone-type=local-zone**
- **machine.openshift.io/zone-group=\$ZONE_GROUP_NAME**

By default, the machine sets for the edge compute pool define the taint of **NoSchedule** to prevent other workloads from spreading on Local Zones instances. Users can only run user workloads if they define tolerations in the pod specification.

Additional resources

- [MTU value selection](#)
- [Changing the MTU for the cluster network](#)
- [Understanding taints and tolerations](#)
- [Storage classes](#)
- [Ingress Controller sharding](#)

6.3.11.3. Installation prerequisites

Before you install a cluster in an AWS Local Zones environment, you must configure your infrastructure so that it can adopt Local Zone capabilities.

6.3.11.3.1. Opting in to an AWS Local Zones

If you plan to create subnets in AWS Local Zones, you must opt in to each zone group separately.

Prerequisites

- You have installed the AWS CLI.
- You have determined an AWS Region for where you want to deploy your OpenShift Container Platform cluster.
- You have attached a permissive IAM policy to a user or role account that opts in to the zone group.

Procedure

1. List the zones that are available in your AWS Region by running the following command:

Example command for listing available AWS Local Zones in an AWS Region

```
$ aws --region "<value_of_AWS_Region>" ec2 describe-availability-zones \
  --query 'AvailabilityZones[].[{ZoneName: ZoneName, GroupName: GroupName, Status:
  OptInStatus}]' \
  --filters Name=zone-type,Values=local-zone \
  --all-availability-zones
```

Depending on the AWS Region, the list of available zones might be long. The command returns the following fields:

ZoneName

The name of the Local Zones.

GroupName

The group that comprises the zone. To opt in to the Region, save the name.

Status

The status of the Local Zones group. If the status is **not-opted-in**, you must opt in the **GroupName** as described in the next step.

2. Opt in to the zone group on your AWS account by running the following command:

```
$ aws ec2 modify-availability-zone-group \
  --group-name "<value_of_GroupName>" 1 \
  --opt-in-status opted-in
```

- 1** Replace **<value_of_GroupName>** with the name of the group of the Local Zones where you want to create subnets. For example, specify **us-east-1-nyc-1** to use the zone **us-east-1-nyc-1a** (US East New York).

6.3.11.3.2. Obtaining an AWS Marketplace image

If you are deploying an OpenShift Container Platform cluster using an AWS Marketplace image, you must first subscribe through AWS. Subscribing to the offer provides you with the AMI ID that the installation program uses to deploy compute nodes.

Prerequisites

- You have an AWS account to purchase the offer. This account does not have to be the same account that is used to install the cluster.

Procedure

1. Complete the OpenShift Container Platform subscription from the [AWS Marketplace](#).
2. Record the AMI ID for your specific AWS Region. As part of the installation process, you must update the **install-config.yaml** file with this value before deploying the cluster.

Sample install-config.yaml file with AWS Marketplace compute nodes

```

apiVersion: v1
baseDomain: example.com
compute:
- hyperthreading: Enabled
  name: worker
  platform:
    aws:
      amiID: ami-06c4d345f7c207239 1
      type: m5.4xlarge
      replicas: 3
  metadata:
    name: test-cluster
  platform:
    aws:
      region: us-east-2 2
  sshKey: ssh-ed25519 AAAA...
  pullSecret: '{"auths": ...}'

```

- 1** The AMI ID from your AWS Marketplace subscription.
- 2** Your AMI ID is associated with a specific AWS Region. When creating the installation configuration file, ensure that you select the same AWS Region that you specified when configuring your subscription.

6.3.11.4. Preparing for the installation

Before you extend nodes to Local Zones, you must prepare certain resources for the cluster installation environment.

6.3.11.4.1. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 6.19. Minimum resource requirements

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Bootstrap	RHCOS	4	16 GB	100 GB	300

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Control plane	RHCOS	4	16 GB	100 GB	300
Compute	RHCOS, RHEL 8.6 and later [3]	2	8 GB	100 GB	300

1. One vCPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio: (threads per core × cores) × sockets = vCPUs.
2. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes which require a 10 ms p99 fsync duration. Note that on many cloud platforms, storage size and IOPS scale together, so you might need to over-allocate storage volume to obtain sufficient performance.
3. As with all user-provisioned installations, if you choose to use RHEL compute machines in your cluster, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Use of RHEL 7 compute machines is deprecated and has been removed in OpenShift Container Platform 4.10 and later.



NOTE

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

6.3.11.4.2. Tested instance types for AWS

The following Amazon Web Services (AWS) instance types have been tested with OpenShift Container Platform for use with AWS Local Zones.

**NOTE**

Use the machine types included in the following charts for your AWS instances. If you use an instance type that is not listed in the chart, ensure that the instance size you use matches the minimum resource requirements that are listed in the section named "Minimum resource requirements for cluster installation".

Example 6.48. Machine types based on 64-bit x86 architecture for AWS Local Zones

- **c5.***
- **c5d.***
- **m6i.***
- **m5.***
- **r5.***
- **t3.***

Additional resources

- See [AWS Local Zones features](#) in the AWS documentation.

6.3.11.4.3. Creating the installation configuration file

Generate and customize the installation configuration file that the installation program needs to deploy your cluster.

Prerequisites

- You obtained the OpenShift Container Platform installation program for user-provisioned infrastructure and the pull secret for your cluster.
- You checked that you are deploying your cluster to an AWS Region with an accompanying Red Hat Enterprise Linux CoreOS (RHCOS) AMI published by Red Hat. If you are deploying to an AWS Region that requires a custom AMI, such as an AWS GovCloud Region, you must create the **install-config.yaml** file manually.

Procedure

1. Create the **install-config.yaml** file.
 - a. Change to the directory that contains the installation program and run the following command:

```
$. ./openshift-install create install-config --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the directory name to store the files that the installation program creates.

**IMPORTANT**

Specify an empty directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

- b. At the prompts, provide the configuration details for your cloud:
 - i. Optional: Select an SSH key to use to access your cluster machines.

**NOTE**


For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- ii. Select **aws** as the platform to target.
- iii. If you do not have an AWS profile stored on your computer, enter the AWS access key ID and secret access key for the user that you configured to run the installation program.

**NOTE**

The AWS access key ID and secret access key are stored in `~/.aws/credentials` in the home directory of the current user on the installation host. You are prompted for the credentials by the installation program if the credentials for the exported profile are not present in the file. Any credentials that you provide to the installation program are stored in the file.

- iv. Select the AWS Region to deploy the cluster to.
 - v. Select the base domain for the Route 53 service that you configured for your cluster.
 - vi. Enter a descriptive name for your cluster.
 - vii. Paste the [pull secret from Red Hat OpenShift Cluster Manager](#) .
2. Optional: Back up the **install-config.yaml** file.

**IMPORTANT**

The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

6.3.11.4.4. Examples of installation configuration files with edge compute pools

The following examples show **install-config.yaml** files that contain an edge machine pool configuration.

Configuration that uses an edge pool with a custom instance type

```
apiVersion: v1
baseDomain: devcluster.openshift.com
metadata:
  name: ipi-edgezone
compute:
- name: edge
  platform:
    aws:
      type: r5.2xlarge
platform:
  aws:
    region: us-west-2
pullSecret: '{"auths": ...}'
sshKey: ssh-ed25519 AAAA...
```

Instance types differ between locations. To verify availability in the Local Zones in which the cluster runs, see the AWS documentation.

Configuration that uses an edge pool with a custom Amazon Elastic Block Store (EBS) type

```
apiVersion: v1
baseDomain: devcluster.openshift.com
metadata:
  name: ipi-edgezone
compute:
- name: edge
  platform:
    aws:
      zones:
      - us-west-2-lax-1a
      - us-west-2-lax-1b
      - us-west-2-phx-2a
      rootVolume:
        type: gp3
        size: 120
platform:
  aws:
    region: us-west-2
pullSecret: '{"auths": ...}'
sshKey: ssh-ed25519 AAAA...
```

Elastic Block Storage (EBS) types differ between locations. Check the AWS documentation to verify availability in the Local Zones in which the cluster runs.

Configuration that uses an edge pool with custom security groups

```
apiVersion: v1
baseDomain: devcluster.openshift.com
metadata:
  name: ipi-edgezone
compute:
- name: edge
  platform:
```

```
aws:
  additionalSecurityGroupIDs:
    - sg-1 1
    - sg-2
platform:
  aws:
    region: us-west-2
pullSecret: '{"auths": ...}'
sshKey: ssh-ed25519 AAAA...
```

- 1 Specify the name of the security group as it is displayed on the Amazon EC2 console. Ensure that you include the **sg** prefix.

6.3.11.4.5. Customizing the cluster network MTU

Before you deploy a cluster on AWS, you can customize the cluster network maximum transmission unit (MTU) for your cluster network to meet the needs of your infrastructure.

By default, when you install a cluster with supported Local Zones capabilities, the MTU value for the cluster network is automatically adjusted to the lowest value that the network plugin accepts.




IMPORTANT

Setting an unsupported MTU value for EC2 instances that operate in the Local Zones infrastructure can cause issues for your OpenShift Container Platform cluster.

If the Local Zone supports higher MTU values in between EC2 instances in the Local Zone and the AWS Region, you can manually configure the higher value to increase the network performance of the cluster network.

You can customize the MTU for a cluster by specifying the **networking.clusterNetworkMTU** parameter in the **install-config.yaml** configuration file.



IMPORTANT

All subnets in Local Zones must support the higher MTU value, so that each node in that zone can successfully communicate with services in the AWS Region and deploy your workloads.

Example of overwriting the default MTU value

```
apiVersion: v1
baseDomain: devcluster.openshift.com
metadata:
  name: edge-zone
networking:
  clusterNetworkMTU: 8901
compute:
  - name: edge
platform:
  aws:
    zones:
      - us-west-2-lax-1a
```

```
- us-west-2-lax-1b
platform:
  aws:
    region: us-west-2
  pullSecret: '{"auths": ...}'
  sshKey: ssh-ed25519 AAAA...
```

Additional resources

- For more information about the maximum supported maximum transmission unit (MTU) value, see [AWS resources supported in Local Zones](#) in the AWS documentation.

6.3.11.5. Cluster installation options for an AWS Local Zones environment

Choose one of the following installation options to install an OpenShift Container Platform cluster on AWS with edge compute nodes defined in Local Zones:

- Fully automated option: Installing a cluster to quickly extend compute nodes to edge compute pools, where the installation program automatically creates infrastructure resources for the OpenShift Container Platform cluster.
- Existing VPC option: Installing a cluster on AWS into an existing VPC, where you supply Local Zones subnets to the **install-config.yaml** file.

Next steps

Choose one of the following options to install an OpenShift Container Platform cluster in an AWS Local Zones environment:

- [Installing a cluster quickly in AWS Local Zones](#)
- [Installing a cluster in an existing VPC with defined AWS Local Zone subnets](#)

6.3.11.6. Install a cluster quickly in AWS Local Zones

For OpenShift Container Platform 4.16, you can quickly install a cluster on Amazon Web Services (AWS) to extend compute nodes to Local Zones locations. By using this installation route, the installation program automatically creates network resources and Local Zones subnets for each zone that you defined in your configuration file. To customize the installation, you must modify parameters in the **install-config.yaml** file before you deploy the cluster.

6.3.11.6.1. Modifying an installation configuration file to use AWS Local Zones

Modify an **install-config.yaml** file to include AWS Local Zones.

Prerequisites

- You have configured an AWS account.
- You added your AWS keys and AWS Region to your local AWS profile by running **aws configure**.
- You are familiar with the configuration limitations that apply when you specify the installation program to automatically create subnets for your OpenShift Container Platform cluster.
- You opted in to the Local Zones group for each zone.

- You created an **install-config.yaml** file by using the procedure "Creating the installation configuration file".

Procedure

1. Modify the **install-config.yaml** file by specifying Local Zones names in the **platform.aws.zones** property of the edge compute pool.

```
# ...
platform:
  aws:
    region: <region_name> ❶
  compute:
  - name: edge
    platform:
      aws:
        zones: ❷
        - <local_zone_name>
#...
```

- ❶ The AWS Region name.
- ❷ The list of Local Zones names that you use must exist in the same AWS Region specified in the **platform.aws.region** field.

Example of a configuration to install a cluster in the `us-west-2` AWS Region that extends edge nodes to Local Zones in Los Angeles and Las Vegas locations

```
apiVersion: v1
baseDomain: example.com
metadata:
  name: cluster-name
platform:
  aws:
    region: us-west-2
  compute:
  - name: edge
    platform:
      aws:
        zones:
        - us-west-2-lax-1a
        - us-west-2-lax-1b
        - us-west-2-las-1a
  pullSecret: '{"auths": ...}'
  sshKey: 'ssh-ed25519 AAAA...'
#...
```

2. Deploy your cluster.

Additional resources

- [Creating the installation configuration file](#)
- [Cluster limitations in AWS Local Zones](#)

Next steps

- [Deploying the cluster](#)

6.3.11.7. Installing a cluster in an existing VPC that has Local Zone subnets

You can install a cluster into an existing Amazon Virtual Private Cloud (VPC) on Amazon Web Services (AWS). The installation program provisions the rest of the required infrastructure, which you can further customize. To customize the installation, modify parameters in the **install-config.yaml** file before you install the cluster.

Installing a cluster on AWS into an existing VPC requires extending compute nodes to the edge of the Cloud Infrastructure by using AWS Local Zones.

Local Zone subnets extend regular compute nodes to edge networks. Each edge compute nodes runs a user workload. After you create an Amazon Web Service (AWS) Local Zone environment, and you deploy your cluster, you can use edge compute nodes to create user workloads in Local Zone subnets.



NOTE

If you want to create private subnets, you must either modify the provided CloudFormation template or create your own template.

You can use a provided CloudFormation template to create network resources. Additionally, you can modify a template to customize your infrastructure or use the information that they contain to create AWS resources according to your company's policies.



IMPORTANT

The steps for performing an installer-provisioned infrastructure installation are provided for example purposes only. Installing a cluster in an existing VPC requires that you have knowledge of the cloud provider and the installation process of OpenShift Container Platform. You can use a CloudFormation template to assist you with completing these steps or to help model your own cluster installation. Instead of using the CloudFormation template to create resources, you can decide to use other methods for generating these resources.

6.3.11.7.1. Creating a VPC in AWS

You can create a Virtual Private Cloud (VPC), and subnets for all Local Zones locations, in Amazon Web Services (AWS) for your OpenShift Container Platform cluster to extend compute nodes to edge locations. You can further customize your VPC to meet your requirements, including a VPN and route tables. You can also add new Local Zones subnets not included at initial deployment.

You can use the provided CloudFormation template and a custom parameter file to create a stack of AWS resources that represent the VPC.



NOTE

If you do not use the provided CloudFormation template to create your AWS infrastructure, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites


- You configured an AWS account.
- You added your AWS keys and AWS Region to your local AWS profile by running **aws configure**.
- You opted in to the AWS Local Zones on your AWS account.

Procedure

1. Create a JSON file that contains the parameter values that the CloudFormation template requires:

```
[
  {
    "ParameterKey": "VpcCidr", 1
    "ParameterValue": "10.0.0.0/16" 2
  },
  {
    "ParameterKey": "AvailabilityZoneCount", 3
    "ParameterValue": "3" 4
  },
  {
    "ParameterKey": "SubnetBits", 5
    "ParameterValue": "12" 6
  }
]
```

- 1** The CIDR block for the VPC.
 - 2** Specify a CIDR block in the format **x.x.x.x/16-24**.
 - 3** The number of availability zones to deploy the VPC in.
 - 4** Specify an integer between **1** and **3**.
 - 5** The size of each subnet in each availability zone.
 - 6** Specify an integer between **5** and **13**, where **5** is **/27** and **13** is **/19**.
2. Go to the section of the documentation named "CloudFormation template for the VPC", and then copy the syntax from the provided template. Save the copied template syntax as a YAML file on your local system. This template describes the VPC that your cluster requires.
 3. Launch the CloudFormation template to create a stack of AWS resources that represent the VPC by running the following command:



IMPORTANT

You must enter the command on a single line.

```
$ aws cloudformation create-stack --stack-name <name> \ 1
--template-body file://<template>.yaml \ 2
--parameters file://<parameters>.json 3
```

- 1 **<name>** is the name for the CloudFormation stack, such as **cluster-vpc**. You need the name of this stack if you remove the cluster.
- 2 **<template>** is the relative path to and name of the CloudFormation template YAML file that you saved.
- 3 **<parameters>** is the relative path and the name of the CloudFormation parameters JSON file.

Example output

```
arn:aws:cloudformation:us-east-1:123456789012:stack/cluster-vpc/dbedae40-2fd3-11eb-820e-12a48460849f
```

4. Confirm that the template components exist by running the following command:

```
$ aws cloudformation describe-stacks --stack-name <name>
```

After the **StackStatus** displays **CREATE_COMPLETE**, the output displays values for the following parameters. You must provide these parameter values to the other CloudFormation templates that you run to create your cluster.

VpcId	The ID of your VPC.
PublicSubnetIds	The IDs of the new public subnets.
PrivateSubnetIds	The IDs of the new private subnets.
PublicRouteTableId	The ID of the new public route table ID.

6.3.11.7.2. CloudFormation template for the VPC

You can use the following CloudFormation template to deploy the VPC that you need for your OpenShift Container Platform cluster.

Example 6.49. CloudFormation template for the VPC

```
AWSTemplateFormatVersion: 2010-09-09
Description: Template for Best Practice VPC with 1-3 AZs

Parameters:
  VpcCidr:
    AllowedPattern: ^((([0-9]|[1-9][0-9]|1[0-9]{2}|2[0-4][0-9]|25[0-5])\.)\.)\{3}([0-9]|[1-9][0-9]|1[0-9]{2}|2[0-4][0-9]|25[0-5])\)(\v(1[6-9]|2[0-4]))$
```


ConstraintDescription: CIDR block parameter must be in the form x.x.x.x/**16-24**.

Default: **10.0.0.0/16**

Description: CIDR block for VPC.

Type: String

AvailabilityZoneCount:

ConstraintDescription: "The number of availability zones. (Min: 1, Max: 3)"

MinValue: **1**

MaxValue: **3**

Default: **1**

Description: "How many AZs to create VPC subnets for. (Min: 1, Max: 3)"

Type: Number

SubnetBits:

ConstraintDescription: CIDR block parameter must be in the form x.x.x.x/**19-27**.

MinValue: **5**

MaxValue: **13**

Default: **12**

Description: "Size of each subnet to create within the availability zones. (Min: 5 = /27, Max: 13 = /19)"

Type: Number

Metadata:

AWS::CloudFormation::Interface:

ParameterGroups:

- Label:

default: "Network Configuration"

Parameters:

- VpcCidr

- SubnetBits

- Label:

default: "Availability Zones"

Parameters:

- AvailabilityZoneCount

ParameterLabels:

AvailabilityZoneCount:

default: "Availability Zone Count"

VpcCidr:

default: "VPC CIDR"

SubnetBits:

default: "Bits Per Subnet"

Conditions:

DoAz3: !Equals [**3**, !Ref AvailabilityZoneCount]

DoAz2: !Or [!Equals [**2**, !Ref AvailabilityZoneCount], Condition: DoAz3]

Resources:

VPC:

Type: "AWS::EC2::VPC"

Properties:

EnableDnsSupport: "true"

EnableDnsHostnames: "true"

CidrBlock: !Ref VpcCidr

PublicSubnet:

Type: "AWS::EC2::Subnet"

Properties:

VpcId: !Ref VPC

CidrBlock: !Select [**0**, !Cidr [!Ref VpcCidr, **6**, !Ref SubnetBits]]

```

    AvailabilityZone: !Select
    - 0
    - Fn::GetAZs: !Ref "AWS::Region"
PublicSubnet2:
  Type: "AWS::EC2::Subnet"
  Condition: DoAz2
  Properties:
    Vpclid: !Ref VPC
    CidrBlock: !Select [1, !Cidr [!Ref VpcCidr, 6, !Ref SubnetBits]]
    AvailabilityZone: !Select
    - 1
    - Fn::GetAZs: !Ref "AWS::Region"
PublicSubnet3:
  Type: "AWS::EC2::Subnet"
  Condition: DoAz3
  Properties:
    Vpclid: !Ref VPC
    CidrBlock: !Select [2, !Cidr [!Ref VpcCidr, 6, !Ref SubnetBits]]
    AvailabilityZone: !Select
    - 2
    - Fn::GetAZs: !Ref "AWS::Region"
InternetGateway:
  Type: "AWS::EC2::InternetGateway"
GatewayToInternet:
  Type: "AWS::EC2::VPCGatewayAttachment"
  Properties:
    Vpclid: !Ref VPC
    InternetGatewayId: !Ref InternetGateway
PublicRouteTable:
  Type: "AWS::EC2::RouteTable"
  Properties:
    Vpclid: !Ref VPC
PublicRoute:
  Type: "AWS::EC2::Route"
  DependsOn: GatewayToInternet
  Properties:
    RouteTableId: !Ref PublicRouteTable
    DestinationCidrBlock: 0.0.0.0/0
    GatewayId: !Ref InternetGateway
PublicSubnetRouteTableAssociation:
  Type: "AWS::EC2::SubnetRouteTableAssociation"
  Properties:
    SubnetId: !Ref PublicSubnet
    RouteTableId: !Ref PublicRouteTable
PublicSubnetRouteTableAssociation2:
  Type: "AWS::EC2::SubnetRouteTableAssociation"
  Condition: DoAz2
  Properties:
    SubnetId: !Ref PublicSubnet2
    RouteTableId: !Ref PublicRouteTable
PublicSubnetRouteTableAssociation3:
  Condition: DoAz3
  Type: "AWS::EC2::SubnetRouteTableAssociation"
  Properties:
    SubnetId: !Ref PublicSubnet3
    RouteTableId: !Ref PublicRouteTable

```

```

PrivateSubnet:
  Type: "AWS::EC2::Subnet"
  Properties:
    VpcId: !Ref VPC
    CidrBlock: !Select [3, !Cidr [!Ref VpcCidr, 6, !Ref SubnetBits]]
    AvailabilityZone: !Select
      - 0
      - Fn::GetAZs: !Ref "AWS::Region"
PrivateRouteTable:
  Type: "AWS::EC2::RouteTable"
  Properties:
    VpcId: !Ref VPC
PrivateSubnetRouteTableAssociation:
  Type: "AWS::EC2::SubnetRouteTableAssociation"
  Properties:
    SubnetId: !Ref PrivateSubnet
    RouteTableId: !Ref PrivateRouteTable
NAT:
  DependsOn:
    - GatewayToInternet
  Type: "AWS::EC2::NatGateway"
  Properties:
    AllocationId:
      "Fn::GetAtt":
        - EIP
        - AllocationId
    SubnetId: !Ref PublicSubnet
EIP:
  Type: "AWS::EC2::EIP"
  Properties:
    Domain: vpc
Route:
  Type: "AWS::EC2::Route"
  Properties:
    RouteTableId:
      Ref: PrivateRouteTable
    DestinationCidrBlock: 0.0.0.0/0
    NatGatewayId:
      Ref: NAT
PrivateSubnet2:
  Type: "AWS::EC2::Subnet"
  Condition: DoAz2
  Properties:
    VpcId: !Ref VPC
    CidrBlock: !Select [4, !Cidr [!Ref VpcCidr, 6, !Ref SubnetBits]]
    AvailabilityZone: !Select
      - 1
      - Fn::GetAZs: !Ref "AWS::Region"
PrivateRouteTable2:
  Type: "AWS::EC2::RouteTable"
  Condition: DoAz2
  Properties:
    VpcId: !Ref VPC
PrivateSubnetRouteTableAssociation2:
  Type: "AWS::EC2::SubnetRouteTableAssociation"
  Condition: DoAz2

```

```
Properties:
  SubnetId: !Ref PrivateSubnet2
  RouteTableId: !Ref PrivateRouteTable2
NAT2:
  DependsOn:
  - GatewayToInternet
  Type: "AWS::EC2::NatGateway"
  Condition: DoAz2
  Properties:
    AllocationId:
      "Fn::GetAtt":
      - EIP2
      - AllocationId
    SubnetId: !Ref PublicSubnet2
EIP2:
  Type: "AWS::EC2::EIP"
  Condition: DoAz2
  Properties:
    Domain: vpc
Route2:
  Type: "AWS::EC2::Route"
  Condition: DoAz2
  Properties:
    RouteTableId:
      Ref: PrivateRouteTable2
    DestinationCidrBlock: 0.0.0.0/0
    NatGatewayId:
      Ref: NAT2
PrivateSubnet3:
  Type: "AWS::EC2::Subnet"
  Condition: DoAz3
  Properties:
    VpcId: !Ref VPC
    CidrBlock: !Select [5, !Cidr [!Ref VpcCidr, 6, !Ref SubnetBits]]
    AvailabilityZone: !Select
    - 2
    - Fn::GetAZs: !Ref "AWS::Region"
PrivateRouteTable3:
  Type: "AWS::EC2::RouteTable"
  Condition: DoAz3
  Properties:
    VpcId: !Ref VPC
PrivateSubnetRouteTableAssociation3:
  Type: "AWS::EC2::SubnetRouteTableAssociation"
  Condition: DoAz3
  Properties:
    SubnetId: !Ref PrivateSubnet3
    RouteTableId: !Ref PrivateRouteTable3
NAT3:
  DependsOn:
  - GatewayToInternet
  Type: "AWS::EC2::NatGateway"
  Condition: DoAz3
  Properties:
    AllocationId:
      "Fn::GetAtt":
```

```

    - EIP3
    - AllocationId
    SubnetId: !Ref PublicSubnet3
EIP3:
  Type: "AWS::EC2::EIP"
  Condition: DoAz3
  Properties:
    Domain: vpc
Route3:
  Type: "AWS::EC2::Route"
  Condition: DoAz3
  Properties:
    RouteTableId:
      Ref: PrivateRouteTable3
    DestinationCidrBlock: 0.0.0.0/0
    NatGatewayId:
      Ref: NAT3
S3Endpoint:
  Type: AWS::EC2::VPCEndpoint
  Properties:
    PolicyDocument:
      Version: 2012-10-17
      Statement:
        - Effect: Allow
          Principal: '*'
          Action:
            - '*'
          Resource:
            - '*'
    RouteTableIds:
      - !Ref PublicRouteTable
      - !Ref PrivateRouteTable
      - !If [DoAz2, !Ref PrivateRouteTable2, !Ref "AWS::NoValue"]
      - !If [DoAz3, !Ref PrivateRouteTable3, !Ref "AWS::NoValue"]
    ServiceName: !Join
      - "
      - - com.amazonaws.
      - - !Ref 'AWS::Region'
      - - .s3
    VpId: !Ref VPC

Outputs:
VpId:
  Description: ID of the new VPC.
  Value: !Ref VPC
PublicSubnetIds:
  Description: Subnet IDs of the public subnets.
  Value:
    !Join [
      " ",
      [!Ref PublicSubnet, !If [DoAz2, !Ref PublicSubnet2, !Ref "AWS::NoValue"], !If [DoAz3, !Ref
PublicSubnet3, !Ref "AWS::NoValue"]]
    ]
PrivateSubnetIds:
  Description: Subnet IDs of the private subnets.
  Value:

```

```

!Join [
  "",
  [!Ref PrivateSubnet, !If [DoAz2, !Ref PrivateSubnet2, !Ref "AWS::NoValue"], !If [DoAz3, !Ref
PrivateSubnet3, !Ref "AWS::NoValue"]]
]
PublicRouteTableId:
  Description: Public Route table ID
  Value: !Ref PublicRouteTable
PrivateRouteTableIds:
  Description: Private Route table IDs
  Value:
    !Join [
      "",
      [
        !Join ["=", [
          !Select [0, "Fn::GetAZs": !Ref "AWS::Region"],
          !Ref PrivateRouteTable
        ]],
        !If [DoAz2,
          !Join ["=", [!Select [1, "Fn::GetAZs": !Ref "AWS::Region"], !Ref PrivateRouteTable2]],
          !Ref "AWS::NoValue"
        ],
        !If [DoAz3,
          !Join ["=", [!Select [2, "Fn::GetAZs": !Ref "AWS::Region"], !Ref PrivateRouteTable3]],
          !Ref "AWS::NoValue"
        ]
      ]
    ]
]

```

6.3.11.7.3. Creating subnets in Local Zones

Before you configure a machine set for edge compute nodes in your OpenShift Container Platform cluster, you must create the subnets in Local Zones. Complete the following procedure for each Local Zone that you want to deploy compute nodes to.

You can use the provided CloudFormation template and create a CloudFormation stack. You can then use this stack to custom provision a subnet.



NOTE

If you do not use the provided CloudFormation template to create your AWS infrastructure, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- You configured an AWS account.
- You added your AWS keys and region to your local AWS profile by running **aws configure**.
- You opted in to the Local Zones group.

Procedure

1. Go to the section of the documentation named "CloudFormation template for the VPC subnet", and copy the syntax from the template. Save the copied template syntax as a YAML file on your local system. This template describes the VPC that your cluster requires.
2. Run the following command to deploy the CloudFormation template, which creates a stack of AWS resources that represent the VPC:

```
$ aws cloudformation create-stack --stack-name <stack_name> \ 1
--region ${CLUSTER_REGION} \
--template-body file://<template>.yaml \ 2
--parameters \
ParameterKey=VpcId,ParameterValue="${VPC_ID}" \ 3
ParameterKey=ClusterName,ParameterValue="${CLUSTER_NAME}" \ 4
ParameterKey=ZoneName,ParameterValue="${ZONE_NAME}" \ 5
ParameterKey=PublicRouteTableId,ParameterValue="${ROUTE_TABLE_PUB}" \ 6
ParameterKey=PublicSubnetCidr,ParameterValue="${SUBNET_CIDR_PUB}" \ 7
ParameterKey=PrivateRouteTableId,ParameterValue="${ROUTE_TABLE_PVT}" \ 8
ParameterKey=PrivateSubnetCidr,ParameterValue="${SUBNET_CIDR_PVT}" \ 9
```

- 1 **<stack_name>** is the name for the CloudFormation stack, such as **cluster-wl-
<local_zone_shortcode>**. You need the name of this stack if you remove the cluster.
- 2 **<template>** is the relative path and the name of the CloudFormation template YAML file that you saved.
- 3 **`\${VPC_ID}`** is the VPC ID, which is the value **VpcId** in the output of the CloudFormation template for the VPC.
- 4 **`\${ZONE_NAME}`** is the value of Local Zones name to create the subnets.
- 5 **`\${CLUSTER_NAME}`** is the value of **ClusterName** to be used as a prefix of the new AWS resource names.
- 6 **`\${SUBNET_CIDR_PUB}`** is a valid CIDR block that is used to create the public subnet. This block must be part of the VPC CIDR block **VpcCidr**.
- 7 **`\${ROUTE_TABLE_PVT}`** is the **PrivateRouteTableId** extracted from the output of the VPC's CloudFormation stack.
- 8 **`\${SUBNET_CIDR_PVT}`** is a valid CIDR block that is used to create the private subnet. This block must be part of the VPC CIDR block **VpcCidr**.

Example output

```
arn:aws:cloudformation:us-east-1:123456789012:stack/<stack_name>/dbedae40-820e-11eb-2fd3-12a48460849f
```

Verification

- Confirm that the template components exist by running the following command:

```
$ aws cloudformation describe-stacks --stack-name <stack_name>
```

After the **StackStatus** displays **CREATE_COMPLETE**, the output displays values for the following parameters. Ensure that you provide these parameter values to the other CloudFormation templates that you run to create for your cluster.

PublicSubnetId	The IDs of the public subnet created by the CloudFormation stack.
PrivateSubnetId	The IDs of the private subnet created by the CloudFormation stack.

6.3.11.7.4. CloudFormation template for the VPC subnet

You can use the following CloudFormation template to deploy the private and public subnets in a zone on Local Zones infrastructure.

Example 6.50. CloudFormation template for VPC subnets

```

AWSTemplateFormatVersion: 2010-09-09
Description: Template for Best Practice Subnets (Public and Private)

Parameters:
  VpcId:
    Description: VPC ID that comprises all the target subnets.
    Type: String
    AllowedPattern: ^(?:(:vpc)(?:-[a-zA-Z0-9]+)?\b(?:[0-9]{1,3}\.){3}[0-9]{1,3})$
    ConstraintDescription: VPC ID must be with valid name, starting with vpc-.*
  ClusterName:
    Description: Cluster name or prefix name to prepend the Name tag for each subnet.
    Type: String
    AllowedPattern: ".+"
    ConstraintDescription: ClusterName parameter must be specified.
  ZoneName:
    Description: Zone Name to create the subnets, such as us-west-2-lax-1a.
    Type: String
    AllowedPattern: ".+"
    ConstraintDescription: ZoneName parameter must be specified.
  PublicRouteTableId:
    Description: Public Route Table ID to associate the public subnet.
    Type: String
    AllowedPattern: ".+"
    ConstraintDescription: PublicRouteTableId parameter must be specified.
  PublicSubnetCidr:
    AllowedPattern: ^(((0-9)[1-9][0-9]|1[0-9]{2}|2[0-4][0-9]|25[0-5])\.)\.{3}((0-9)[1-9][0-9]|1[0-9]{2}|2[0-4][0-9]|25[0-5])(\/(1[6-9]|2[0-4]))$
    ConstraintDescription: CIDR block parameter must be in the form x.x.x.x/16-24.
    Default: 10.0.128.0/20
    Description: CIDR block for public subnet.
    Type: String
  PrivateRouteTableId:
    Description: Private Route Table ID to associate the private subnet.
    Type: String
    AllowedPattern: ".+"
    ConstraintDescription: PrivateRouteTableId parameter must be specified.
  PrivateSubnetCidr:

```



```

AllowedPattern: ^(((0-9)[1-9][0-9]|1[0-9]{2}|2[0-4][0-9]|25[0-5])\.)\{3\}((0-9)[1-9][0-9]|1[0-9]{2}|2[0-4][0-9]|25[0-5])\(\v(1[6-9]|2[0-4])\))$
ConstraintDescription: CIDR block parameter must be in the form x.x.x.x/16-24.
Default: 10.0.128.0/20
Description: CIDR block for private subnet.
Type: String

```

Resources:

PublicSubnet:

Type: "AWS::EC2::Subnet"

Properties:

VpcId: !Ref VpcId

CidrBlock: !Ref PublicSubnetCidr

AvailabilityZone: !Ref ZoneName

Tags:

- Key: Name

Value: !Join ['-', [!Ref ClusterName, "public", !Ref ZoneName]]

PublicSubnetRouteTableAssociation:

Type: "AWS::EC2::SubnetRouteTableAssociation"

Properties:

SubnetId: !Ref PublicSubnet

RouteTableId: !Ref PublicRouteTableId

PrivateSubnet:

Type: "AWS::EC2::Subnet"

Properties:

VpcId: !Ref VpcId

CidrBlock: !Ref PrivateSubnetCidr

AvailabilityZone: !Ref ZoneName

Tags:

- Key: Name

Value: !Join ['-', [!Ref ClusterName, "private", !Ref ZoneName]]

PrivateSubnetRouteTableAssociation:

Type: "AWS::EC2::SubnetRouteTableAssociation"

Properties:

SubnetId: !Ref PrivateSubnet

RouteTableId: !Ref PrivateRouteTableId

Outputs:

PublicSubnetId:

Description: Subnet ID of the public subnets.

Value:

!Join [""], [!Ref PublicSubnet]]

PrivateSubnetId:

Description: Subnet ID of the private subnets.

Value:

!Join [""], [!Ref PrivateSubnet]]

Additional resources

- You can view details about the CloudFormation stacks that you create by navigating to the [AWS CloudFormation console](#).

6.3.11.7.5. Modifying an installation configuration file to use AWS Local Zones subnets

Modify your **install-config.yaml** file to include Local Zones subnets.

Prerequisites

- You created subnets by using the procedure "Creating subnets in Local Zones".
- You created an **install-config.yaml** file by using the procedure "Creating the installation configuration file".

Procedure

- Modify the **install-config.yaml** configuration file by specifying Local Zones subnets in the **platform.aws.subnets** parameter.

Example installation configuration file with Local Zones subnets

```
# ...
platform:
  aws:
    region: us-west-2
    subnets: 1
      - publicSubnetId-1
      - publicSubnetId-2
      - publicSubnetId-3
      - privateSubnetId-1
      - privateSubnetId-2
      - privateSubnetId-3
      - publicSubnetId-LocalZone-1
# ...
```

- 1** List of subnet IDs created in the zones: Availability and Local Zones.

Additional resources

- For more information about viewing the CloudFormation stacks that you created, see [AWS CloudFormation console](#).
- For more information about AWS profile and credential configuration, see [Configuration and credential file settings](#) in the AWS documentation.

Next steps

- [Deploying the cluster](#)

6.3.11.8. Optional: AWS security groups

By default, the installation program creates and attaches security groups to control plane and compute machines. The rules associated with the default security groups cannot be modified.

However, you can apply additional existing AWS security groups, which are associated with your existing VPC, to control plane and compute machines. Applying custom security groups can help you meet the security needs of your organization, in such cases where you need to control the incoming or outgoing traffic of these machines.

As part of the installation process, you apply custom security groups by modifying the **install-config.yaml** file before deploying the cluster.


For more information, see "Edge compute pools and AWS Local Zones".

6.3.11.9. Optional: Assign public IP addresses to edge compute nodes

If your workload requires deploying the edge compute nodes in public subnets on Local Zones infrastructure, you can configure the machine set manifests when installing a cluster.

AWS Local Zones infrastructure accesses the network traffic in a specified zone, so applications can take advantage of lower latency when serving end users that are closer to that zone.

The default setting that deploys compute nodes in private subnets might not meet your needs, so consider creating edge compute nodes in public subnets when you want to apply more customization to your infrastructure.



IMPORTANT

By default, OpenShift Container Platform deploys the compute nodes in private subnets. For best performance, consider placing compute nodes in subnets that have their Public IP addresses attached to the subnets.

You must create additional security groups, but ensure that you only open the groups' rules over the internet when you really need to.

Procedure

1. Change to the directory that contains the installation program and generate the manifest files. Ensure that the installation manifests get created at the **openshift** and **manifests** directory level.

```
$ ./openshift-install create manifests --dir <installation_directory>
```

2. Edit the machine set manifest that the installation program generates for the Local Zones, so that the manifest gets deployed in public subnets. Specify **true** for the **spec.template.spec.providerSpec.value.publicIP** parameter.

Example machine set manifest configuration for installing a cluster quickly in Local Zones

```
spec:
  template:
    spec:
      providerSpec:
        value:
          publicIP: true
          subnet:
            filters:
```

```
- name: tag:Name
  values:
    - ${INFRA_ID}-public-${ZONE_NAME}
```

Example machine set manifest configuration for installing a cluster in an existing VPC that has Local Zones subnets

```
apiVersion: machine.openshift.io/v1beta1
kind: MachineSet
metadata:
  name: <infrastructure_id>-edge-<zone>
  namespace: openshift-machine-api
spec:
  template:
    spec:
      providerSpec:
        value:
          publicIp: true
```

6.3.11.10. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.


Procedure

1. Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
  --log-level=info 2
```

- 1** For **<installation_directory>**, specify the location of your customized **./install-config.yaml** file.
- 2** To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

- Optional: Remove or disable the **AdministratorAccess** policy from the IAM account that you used to install the cluster.




NOTE

The elevated permissions provided by the **AdministratorAccess** policy are required only during installation.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to `<installation_directory>/openshift_install.log`.




IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

6.3.11.11. Verifying the status of the deployed cluster

Verify that your OpenShift Container Platform successfully deployed on AWS Local Zones.

6.3.11.1.1. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

6.3.11.1.2. Logging in to the cluster by using the web console

The **kubeadmin** user exists by default after an OpenShift Container Platform installation. You can log in to your cluster as the **kubeadmin** user by using the OpenShift Container Platform web console.

Prerequisites

- You have access to the installation host.
- You completed a cluster installation and all cluster Operators are available.

Procedure

1. Obtain the password for the **kubeadmin** user from the **kubeadmin-password** file on the installation host:

```
$ cat <installation_directory>/auth/kubeadmin-password
```



NOTE

Alternatively, you can obtain the **kubeadmin** password from the **<installation_directory>/openshift_install.log** log file on the installation host.

- List the OpenShift Container Platform web console route:

```
$ oc get routes -n openshift-console | grep 'console-openshift'
```



NOTE

Alternatively, you can obtain the OpenShift Container Platform route from the `<installation_directory>/openshift_install.log` log file on the installation host.

Example output

```
console    console-openshift-console.apps.<cluster_name>.<base_domain>    console
https    reencrypt/Redirect    None
```

- Navigate to the route detailed in the output of the preceding command in a web browser and log in as the **kubeadmin** user.

Additional resources

- [Accessing the web console](#)

6.3.11.11.3. Verifying nodes that were created with edge compute pool

After you install a cluster that uses AWS Local Zones infrastructure, check the status of the machine that was created by the machine set manifests created during installation.

- To check the machine sets created from the subnet you added to the **install-config.yaml** file, run the following command:

```
$ oc get machineset -n openshift-machine-api
```

Example output

NAME	DESIRED	CURRENT	READY	AVAILABLE	AGE
cluster-7xw5g-edge-us-east-1-nyc-1a	1	1	1	1	3h4m
cluster-7xw5g-worker-us-east-1a	1	1	1	1	3h4m
cluster-7xw5g-worker-us-east-1b	1	1	1	1	3h4m
cluster-7xw5g-worker-us-east-1c	1	1	1	1	3h4m

- To check the machines that were created from the machine sets, run the following command:

```
$ oc get machines -n openshift-machine-api
```

Example output

NAME	PHASE	TYPE	REGION	ZONE	AGE
cluster-7xw5g-edge-us-east-1-nyc-1a-wbclh-nyc-1a	Running	m6i.xlarge	c5d.2xlarge	us-east-1	us-east-1-3h
cluster-7xw5g-master-0	Running	m6i.xlarge	us-east-1	us-east-1a	3h4m
cluster-7xw5g-master-1	Running	m6i.xlarge	us-east-1	us-east-1b	3h4m
cluster-7xw5g-master-2	Running	m6i.xlarge	us-east-1	us-east-1c	3h4m
cluster-7xw5g-worker-us-east-1a-rtp45	Running	m6i.xlarge	us-east-1	us-east-1a	

```

3h
cluster-7xw5g-worker-us-east-1b-glm7c    Running  m6i.xlarge  us-east-1  us-east-1b
3h
cluster-7xw5g-worker-us-east-1c-qfvz4    Running  m6i.xlarge  us-east-1  us-east-1c
3h

```

- To check nodes with edge roles, run the following command:

```
$ oc get nodes -l node-role.kubernetes.io/edge
```

Example output

```

NAME                                STATUS  ROLES    AGE  VERSION
ip-10-0-207-188.ec2.internal  Ready  edge,worker  172m  v1.25.2+d2e245f

```

Next steps

- [Validating an installation](#).
- If necessary, you can [opt out of remote health](#).

6.3.12. Installing a cluster with compute nodes on AWS Wavelength Zones

You can quickly install an OpenShift Container Platform cluster on Amazon Web Services (AWS) Wavelength Zones by setting the zone names in the edge compute pool of the **install-config.yaml** file, or install a cluster in an existing Amazon Virtual Private Cloud (VPC) with Wavelength Zone subnets.

AWS Wavelength Zones is an infrastructure that AWS configured for mobile edge computing (MEC) applications.

A Wavelength Zone embeds AWS compute and storage services within the 5G network of a communication service provider (CSP). By placing application servers in a Wavelength Zone, the application traffic from your 5G devices can stay in the 5G network. The application traffic of the device reaches the target server directly, making latency a non-issue.

Additional resources

- See [Wavelength Zones](#) in the AWS documentation.

6.3.12.1. Infrastructure prerequisites

- You reviewed details about [OpenShift Container Platform installation and update](#) processes.
- You are familiar with [Selecting a cluster installation method and preparing it for users](#).
- You [configured an AWS account](#) to host the cluster.

**WARNING**

If you have an AWS profile stored on your computer, it must not use a temporary session token that you generated while using a multi-factor authentication device. The cluster continues to use your current AWS credentials to create AWS resources for the entire life of the cluster, so you must use key-based, long-term credentials. To generate appropriate keys, see [Managing Access Keys for IAM Users](#) in the AWS documentation. You can supply the keys when you run the installation program.

- You downloaded the AWS CLI and installed it on your computer. See [Install the AWS CLI Using the Bundled Installer \(Linux, macOS, or UNIX\)](#) in the AWS documentation.
- If you use a firewall, you [configured it to allow the sites](#) that your cluster must access.
- You noted the region and supported [AWS Wavelength Zone locations](#) to create the network resources in.
- You read [AWS Wavelength features](#) in the AWS documentation.
- You read the [Quotas and considerations for Wavelength Zones](#) in the AWS documentation.
- You added permissions for creating network resources that support AWS Wavelength Zones to the Identity and Access Management (IAM) user or role. For example:

Example of an additional IAM policy that attached `ec2:ModifyAvailabilityZoneGroup`, `ec2:CreateCarrierGateway`, and `ec2>DeleteCarrierGateway` permissions to a user or role

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:DeleteCarrierGateway",
        "ec2:CreateCarrierGateway"
      ],
      "Resource": "*"
    },
    {
      "Action": [
        "ec2:ModifyAvailabilityZoneGroup"
      ],
      "Effect": "Allow",
      "Resource": "*"
    }
  ]
}
```

6.3.12.2. About AWS Wavelength Zones and edge compute pool

Read the following sections to understand infrastructure behaviors and cluster limitations in an AWS Wavelength Zones environment.

6.3.12.2.1. Cluster limitations in AWS Wavelength Zones

Some limitations exist when you try to deploy a cluster with a default installation configuration in an Amazon Web Services (AWS) Wavelength Zone.



IMPORTANT

The following list details limitations when deploying a cluster in a pre-configured AWS zone:

- The maximum transmission unit (MTU) between an Amazon EC2 instance in a zone and an Amazon EC2 instance in the Region is **1300**. This causes the cluster-wide network MTU to change according to the network plugin that is used with the deployment.
- Network resources such as Network Load Balancer (NLB), Classic Load Balancer, and Network Address Translation (NAT) Gateways are not globally supported.
- For an OpenShift Container Platform cluster on AWS, the AWS Elastic Block Storage (EBS) **gp3** type volume is the default for node volumes and the default for the storage class. This volume type is not globally available on zone locations. By default, the nodes running in zones are deployed with the **gp2** EBS volume. The **gp2-csi StorageClass** parameter must be set when creating workloads on zone nodes.

If you want the installation program to automatically create Wavelength Zone subnets for your OpenShift Container Platform cluster, specific configuration limitations apply with this method. The following note details some of these limitations. For other limitations, ensure that you read the "Quotas and considerations for Wavelength Zones" document that Red Hat provides in the "Infrastructure prerequisites" section.



IMPORTANT

The following configuration limitation applies when you set the installation program to automatically create subnets for your OpenShift Container Platform cluster:

- When the installation program creates private subnets in AWS Wavelength Zones, the program associates each subnet with the route table of its parent zone. This operation ensures that each private subnet can route egress traffic to the internet by way of NAT Gateways in an AWS Region.
- If the parent-zone route table does not exist during cluster installation, the installation program associates any private subnet with the first available private route table in the Amazon Virtual Private Cloud (VPC). This approach is valid only for AWS Wavelength Zones subnets in an OpenShift Container Platform cluster.

6.3.12.2.2. About edge compute pools

Edge compute nodes are tainted compute nodes that run in AWS Wavelength Zones locations.

When deploying a cluster that uses Wavelength Zones, consider the following points:

- Amazon EC2 instances in the Wavelength Zones are more expensive than Amazon EC2 instances in the Availability Zones.
- The latency is lower between the applications running in AWS Wavelength Zones and the end user. A latency impact exists for some workloads if, for example, ingress traffic is mixed between Wavelength Zones and Availability Zones.



IMPORTANT

Generally, the maximum transmission unit (MTU) between an Amazon EC2 instance in a Wavelength Zones and an Amazon EC2 instance in the Region is 1300. The cluster network MTU must be always less than the EC2 MTU to account for the overhead. The specific overhead is determined by the network plugin. For example: OVN-Kubernetes has an overhead of **100 bytes**.

The network plugin can provide additional features, such as IPsec, that also affect the MTU sizing.

For more information, see [How AWS Wavelength work](#) in the AWS documentation.

OpenShift Container Platform 4.12 introduced a new compute pool, *edge*, that is designed for use in remote zones. The edge compute pool configuration is common between AWS Wavelength Zones locations. Because of the type and size limitations of resources like EC2 and EBS on Wavelength Zones resources, the default instance type can vary from the traditional compute pool.

The default Elastic Block Store (EBS) for Wavelength Zones locations is **gp2**, which differs from the non-edge compute pool. The instance type used for each Wavelength Zones on an edge compute pool also might differ from other compute pools, depending on the instance offerings on the zone.

The edge compute pool creates new labels that developers can use to deploy applications onto AWS Wavelength Zones nodes. The new labels are:

- **node-role.kubernetes.io/edge=""**
- **machine.openshift.io/zone-type=wavelength-zone**
- **machine.openshift.io/zone-group=\$ZONE_GROUP_NAME**

By default, the machine sets for the edge compute pool define the taint of **NoSchedule** to prevent other workloads from spreading on Wavelength Zones instances. Users can only run user workloads if they define tolerations in the pod specification.

Additional resources

- [MTU value selection](#)
- [Changing the MTU for the cluster network](#)
- [Understanding taints and tolerations](#)
- [Storage classes](#)
- [Ingress Controller sharding](#)

6.3.12.3. Installation prerequisites

Before you install a cluster in an AWS Wavelength Zones environment, you must configure your infrastructure so that it can adopt Wavelength Zone capabilities.

6.3.12.3.1. Opting in to an AWS Wavelength Zones

If you plan to create subnets in AWS Wavelength Zones, you must opt in to each zone group separately.

Prerequisites

- You have installed the AWS CLI.
- You have determined an AWS Region for where you want to deploy your OpenShift Container Platform cluster.
- You have attached a permissive IAM policy to a user or role account that opts in to the zone group.

Procedure

1. List the zones that are available in your AWS Region by running the following command:

Example command for listing available AWS Wavelength Zones in an AWS Region

```
$ aws --region "<value_of_AWS_Region>" ec2 describe-availability-zones \
  --query 'AvailabilityZones[].[{ZoneName: ZoneName, GroupName: GroupName, Status:
  OptInStatus}]' \
  --filters Name=zone-type,Values=wavelength-zone \
  --all-availability-zones
```

Depending on the AWS Region, the list of available zones might be long. The command returns the following fields:

ZoneName

The name of the Wavelength Zones.

GroupName

The group that comprises the zone. To opt in to the Region, save the name.

Status

The status of the Wavelength Zones group. If the status is **not-opted-in**, you must opt in the **GroupName** as described in the next step.

2. Opt in to the zone group on your AWS account by running the following command:

```
$ aws ec2 modify-availability-zone-group \
  --group-name "<value_of_GroupName>" 1 \
  --opt-in-status opted-in
```

- 1** Replace **<value_of_GroupName>** with the name of the group of the Wavelength Zones where you want to create subnets. As an example for Wavelength Zones, specify **us-east-1-wl1** to use the zone **us-east-1-wl1-nyc-wlz-1** (US East New York).

6.3.12.3.2. Obtaining an AWS Marketplace image

If you are deploying an OpenShift Container Platform cluster using an AWS Marketplace image, you must first subscribe through AWS. Subscribing to the offer provides you with the AMI ID that the installation program uses to deploy compute nodes.

Prerequisites

- You have an AWS account to purchase the offer. This account does not have to be the same account that is used to install the cluster.

Procedure

1. Complete the OpenShift Container Platform subscription from the [AWS Marketplace](#).
2. Record the AMI ID for your specific AWS Region. As part of the installation process, you must update the **install-config.yaml** file with this value before deploying the cluster.

Sample **install-config.yaml** file with AWS Marketplace compute nodes

```
apiVersion: v1
baseDomain: example.com
compute:
- hyperthreading: Enabled
  name: worker
  platform:
    aws:
      amiID: ami-06c4d345f7c207239 1
      type: m5.4xlarge
    replicas: 3
  metadata:
    name: test-cluster
  platform:
    aws:
      region: us-east-2 2
  sshKey: ssh-ed25519 AAAA...
  pullSecret: '{"auths": ...}'
```

- 1** The AMI ID from your AWS Marketplace subscription.
- 2** Your AMI ID is associated with a specific AWS Region. When creating the installation configuration file, ensure that you select the same AWS Region that you specified when configuring your subscription.

6.3.12.4. Preparing for the installation

Before you extend nodes to Wavelength Zones, you must prepare certain resources for the cluster installation environment.

6.3.12.4.1. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 6.20. Minimum resource requirements

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Bootstrap	RHCOS	4	16 GB	100 GB	300
Control plane	RHCOS	4	16 GB	100 GB	300
Compute	RHCOS, RHEL 8.6 and later [3]	2	8 GB	100 GB	300

1. One vCPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio: (threads per core × cores) × sockets = vCPUs.
2. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes which require a 10 ms p99 fsync duration. Note that on many cloud platforms, storage size and IOPS scale together, so you might need to over-allocate storage volume to obtain sufficient performance.
3. As with all user-provisioned installations, if you choose to use RHEL compute machines in your cluster, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Use of RHEL 7 compute machines is deprecated and has been removed in OpenShift Container Platform 4.10 and later.



NOTE

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

6.3.12.4.2. Tested instance types for AWS

The following Amazon Web Services (AWS) instance types have been tested with OpenShift Container Platform for use with AWS Wavelength Zones.

**NOTE**

Use the machine types included in the following charts for your AWS instances. If you use an instance type that is not listed in the chart, ensure that the instance size you use matches the minimum resource requirements that are listed in the section named "Minimum resource requirements for cluster installation".

Example 6.51. Machine types based on 64-bit x86 architecture for AWS Wavelength Zones

- **r5.***
- **t3.***

Additional resources

- See [AWS Wavelength features](#) in the AWS documentation.

6.3.12.4.3. Creating the installation configuration file

Generate and customize the installation configuration file that the installation program needs to deploy your cluster.

Prerequisites

- You obtained the OpenShift Container Platform installation program and the pull secret for your cluster.
- You checked that you are deploying your cluster to an AWS Region with an accompanying Red Hat Enterprise Linux CoreOS (RHCOS) AMI published by Red Hat. If you are deploying to an AWS Region that requires a custom AMI, such as an AWS GovCloud Region, you must create the **install-config.yaml** file manually.

Procedure

1. Create the **install-config.yaml** file.
 - a. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install create install-config --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the directory name to store the files that the installation program creates.

**IMPORTANT**

Specify an empty directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

- b. At the prompts, provide the configuration details for your cloud:
 - i. Optional: Select an SSH key to use to access your cluster machines.

**NOTE**

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- ii. Select **aws** as the platform to target.
- iii. If you do not have an AWS profile stored on your computer, enter the AWS access key ID and secret access key for the user that you configured to run the installation program.

**NOTE**

The AWS access key ID and secret access key are stored in `~/.aws/credentials` in the home directory of the current user on the installation host. You are prompted for the credentials by the installation program if the credentials for the exported profile are not present in the file. Any credentials that you provide to the installation program are stored in the file.

- iv. Select the AWS Region to deploy the cluster to.
 - v. Select the base domain for the Route 53 service that you configured for your cluster.
 - vi. Enter a descriptive name for your cluster.
 - vii. Paste the [pull secret from Red Hat OpenShift Cluster Manager](#) .
2. Optional: Back up the **install-config.yaml** file.

**IMPORTANT**

The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

6.3.12.4.4. Examples of installation configuration files with edge compute pools

The following examples show **install-config.yaml** files that contain an edge machine pool configuration.

Configuration that uses an edge pool with a custom instance type

```

apiVersion: v1
baseDomain: devcluster.openshift.com
metadata:
  name: ipi-edgezone
compute:
- name: edge
  platform:
    aws:
      type: r5.2xlarge
platform:
  aws:
    region: us-west-2
pullSecret: '{"auths": ...}'
sshKey: ssh-ed25519 AAAA...

```

Instance types differ between locations. To verify availability in the Wavelength Zones in which the cluster runs, see the AWS documentation.

Configuration that uses an edge pool with custom security groups

```

apiVersion: v1
baseDomain: devcluster.openshift.com
metadata:
  name: ipi-edgezone
compute:
- name: edge
  platform:
    aws:
      additionalSecurityGroupIDs:
        - sg-1 1
        - sg-2
platform:
  aws:
    region: us-west-2
pullSecret: '{"auths": ...}'
sshKey: ssh-ed25519 AAAA...

```

- 1** Specify the name of the security group as it is displayed on the Amazon EC2 console. Ensure that you include the **sg** prefix.

6.3.12.5. Cluster installation options for an AWS Wavelength Zones environment

Choose one of the following installation options to install an OpenShift Container Platform cluster on AWS with edge compute nodes defined in Wavelength Zones:

- Fully automated option: Installing a cluster to quickly extend compute nodes to edge compute pools, where the installation program automatically creates infrastructure resources for the OpenShift Container Platform cluster.
- Existing VPC option: Installing a cluster on AWS into an existing VPC, where you supply Wavelength Zones subnets to the **install-config.yaml** file.

Next steps

Choose one of the following options to install an OpenShift Container Platform cluster in an AWS Wavelength Zones environment:

- [Installing a cluster quickly in AWS Wavelength Zones](#)
- [Modifying an installation configuration file to use AWS Wavelength Zones](#)

6.3.12.6. Install a cluster quickly in AWS Wavelength Zones

For OpenShift Container Platform 4.16, you can quickly install a cluster on Amazon Web Services (AWS) to extend compute nodes to Wavelength Zones locations. By using this installation route, the installation program automatically creates network resources and Wavelength Zones subnets for each zone that you defined in your configuration file. To customize the installation, you must modify parameters in the **install-config.yaml** file before you deploy the cluster.

6.3.12.6.1. Modifying an installation configuration file to use AWS Wavelength Zones

Modify an **install-config.yaml** file to include AWS Wavelength Zones.

Prerequisites

- You have configured an AWS account.
- You added your AWS keys and AWS Region to your local AWS profile by running **aws configure**.
- You are familiar with the configuration limitations that apply when you specify the installation program to automatically create subnets for your OpenShift Container Platform cluster.
- You opted in to the Wavelength Zones group for each zone.
- You created an **install-config.yaml** file by using the procedure "Creating the installation configuration file".

Procedure

1. Modify the **install-config.yaml** file by specifying Wavelength Zones names in the **platform.aws.zones** property of the edge compute pool.

```
# ...
platform:
  aws:
    region: <region_name> 1
  compute:
  - name: edge
    platform:
      aws:
        zones: 2
        - <wavelength_zone_name>
#...
```

- 1 The AWS Region name.

- 2 The list of Wavelength Zones names that you use must exist in the same AWS Region specified in the **platform.aws.region** field.

Example of a configuration to install a cluster in the `us-west-2` AWS Region that extends edge nodes to Wavelength Zones in Los Angeles and Las Vegas locations

```
apiVersion: v1
baseDomain: example.com
metadata:
  name: cluster-name
platform:
  aws:
    region: us-west-2
compute:
  - name: edge
    platform:
      aws:
        zones:
          - us-west-2-wl1-lax-wlz-1
          - us-west-2-wl1-las-wlz-1
pullSecret: '{"auths": ...}'
sshKey: 'ssh-ed25519 AAAA...'
#...
```

2. Deploy your cluster.

Additional resources

- [Creating the installation configuration file](#)
- [Cluster limitations in AWS Wavelength Zones](#)

Next steps


- [Deploying the cluster](#)

6.3.12.7. Installing a cluster in an existing VPC that has Wavelength Zone subnets

You can install a cluster into an existing Amazon Virtual Private Cloud (VPC) on Amazon Web Services (AWS). The installation program provisions the rest of the required infrastructure, which you can further customize. To customize the installation, modify parameters in the **install-config.yaml** file before you install the cluster.

Installing a cluster on AWS into an existing VPC requires extending compute nodes to the edge of the Cloud Infrastructure by using AWS Wavelength Zones.

You can use a provided CloudFormation template to create network resources. Additionally, you can modify a template to customize your infrastructure or use the information that they contain to create AWS resources according to your company's policies.



IMPORTANT

The steps for performing an installer-provisioned infrastructure installation are provided for example purposes only. Installing a cluster in an existing VPC requires that you have knowledge of the cloud provider and the installation process of OpenShift Container Platform. You can use a CloudFormation template to assist you with completing these steps or to help model your own cluster installation. Instead of using the CloudFormation template to create resources, you can decide to use other methods for generating these resources.

6.3.12.7.1. Creating a VPC in AWS

You can create a Virtual Private Cloud (VPC), and subnets for all Wavelength Zones locations, in Amazon Web Services (AWS) for your OpenShift Container Platform cluster to extend compute nodes to edge locations. You can further customize your VPC to meet your requirements, including a VPN and route tables. You can also add new Wavelength Zones subnets not included at initial deployment.

You can use the provided CloudFormation template and a custom parameter file to create a stack of AWS resources that represent the VPC.



NOTE

If you do not use the provided CloudFormation template to create your AWS infrastructure, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- You configured an AWS account.
- You added your AWS keys and AWS Region to your local AWS profile by running **aws configure**.
- You opted in to the AWS Wavelength Zones on your AWS account.


Procedure

1. Create a JSON file that contains the parameter values that the CloudFormation template requires:

```
[
  {
    "ParameterKey": "VpcCidr", 1
    "ParameterValue": "10.0.0.0/16" 2
  },
  {
    "ParameterKey": "AvailabilityZoneCount", 3
    "ParameterValue": "3" 4
  },
  {
    "ParameterKey": "SubnetBits", 5
  }
]
```

```
"ParameterValue": "12" 6
}
]
```

- 1 The CIDR block for the VPC.
 - 2 Specify a CIDR block in the format **x.x.x.x/16-24**.
 - 3 The number of availability zones to deploy the VPC in.
 - 4 Specify an integer between **1** and **3**.
 - 5 The size of each subnet in each availability zone.
 - 6 Specify an integer between **5** and **13**, where **5** is **/27** and **13** is **/19**.
2. Go to the section of the documentation named "CloudFormation template for the VPC", and then copy the syntax from the provided template. Save the copied template syntax as a YAML file on your local system. This template describes the VPC that your cluster requires.
 3. Launch the CloudFormation template to create a stack of AWS resources that represent the VPC by running the following command:



IMPORTANT

You must enter the command on a single line.

```
$ aws cloudformation create-stack --stack-name <name> \ 1
--template-body file://<template>.yaml \ 2
--parameters file://<parameters>.json 3
```

- 1 **<name>** is the name for the CloudFormation stack, such as **cluster-vpc**. You need the name of this stack if you remove the cluster.
- 2 **<template>** is the relative path to and name of the CloudFormation template YAML file that you saved.
- 3 **<parameters>** is the relative path and the name of the CloudFormation parameters JSON file.

Example output

```
arn:aws:cloudformation:us-east-1:123456789012:stack/cluster-vpc/dbedae40-2fd3-11eb-
820e-12a48460849f
```

4. Confirm that the template components exist by running the following command:

```
$ aws cloudformation describe-stacks --stack-name <name>
```

After the **StackStatus** displays **CREATE_COMPLETE**, the output displays values for the following parameters. You must provide these parameter values to the other CloudFormation templates that you run to create your cluster.

VpcId	The ID of your VPC.
PublicSubnetIds	The IDs of the new public subnets.
PrivateSubnetIds	The IDs of the new private subnets.
PublicRouteTableId	The ID of the new public route table ID.

6.3.12.7.2. CloudFormation template for the VPC

You can use the following CloudFormation template to deploy the VPC that you need for your OpenShift Container Platform cluster.

Example 6.52. CloudFormation template for the VPC

```

AWSTemplateFormatVersion: 2010-09-09
Description: Template for Best Practice VPC with 1-3 AZs

Parameters:
  VpcCidr:
    AllowedPattern: ^(([0-9]|[1-9][0-9]|1[0-9]{2}|2[0-4][0-9]|25[0-5])\.)\{3\}([0-9]|[1-9][0-9]|1[0-9]{2}|2[0-4][0-9]|25[0-5])(\/(1[6-9]|2[0-4]))$
    ConstraintDescription: CIDR block parameter must be in the form x.x.x.x/16-24.
    Default: 10.0.0.0/16
    Description: CIDR block for VPC.
    Type: String
  AvailabilityZoneCount:
    ConstraintDescription: "The number of availability zones. (Min: 1, Max: 3)"
    MinValue: 1
    MaxValue: 3
    Default: 1
    Description: "How many AZs to create VPC subnets for. (Min: 1, Max: 3)"
    Type: Number
  SubnetBits:
    ConstraintDescription: CIDR block parameter must be in the form x.x.x.x/19-27.
    MinValue: 5
    MaxValue: 13
    Default: 12
    Description: "Size of each subnet to create within the availability zones. (Min: 5 = /27, Max: 13 = /19)"
    Type: Number

Metadata:
  AWS::CloudFormation::Interface:
    ParameterGroups:
      - Label:
          default: "Network Configuration"
        Parameters:
          - VpcCidr
          - SubnetBits

```

```

- Label:
  default: "Availability Zones"
Parameters:
- AvailabilityZoneCount
ParameterLabels:
AvailabilityZoneCount:
  default: "Availability Zone Count"
VpcCidr:
  default: "VPC CIDR"
SubnetBits:
  default: "Bits Per Subnet"

```

Conditions:

```

DoAz3: !Equals [3, !Ref AvailabilityZoneCount]
DoAz2: !Or [!Equals [2, !Ref AvailabilityZoneCount], Condition: DoAz3]

```

Resources:

VPC:

```

Type: "AWS::EC2::VPC"
Properties:
  EnableDnsSupport: "true"
  EnableDnsHostnames: "true"
  CidrBlock: !Ref VpcCidr

```

PublicSubnet:

```

Type: "AWS::EC2::Subnet"
Properties:
  VpcId: !Ref VPC
  CidrBlock: !Select [0, !Cidr [!Ref VpcCidr, 6, !Ref SubnetBits]]
  AvailabilityZone: !Select
  - 0
  - Fn::GetAZs: !Ref "AWS::Region"

```

PublicSubnet2:

```

Type: "AWS::EC2::Subnet"
Condition: DoAz2
Properties:
  VpcId: !Ref VPC
  CidrBlock: !Select [1, !Cidr [!Ref VpcCidr, 6, !Ref SubnetBits]]
  AvailabilityZone: !Select
  - 1
  - Fn::GetAZs: !Ref "AWS::Region"

```

PublicSubnet3:

```

Type: "AWS::EC2::Subnet"
Condition: DoAz3
Properties:
  VpcId: !Ref VPC
  CidrBlock: !Select [2, !Cidr [!Ref VpcCidr, 6, !Ref SubnetBits]]
  AvailabilityZone: !Select
  - 2
  - Fn::GetAZs: !Ref "AWS::Region"

```

InternetGateway:

```

Type: "AWS::EC2::InternetGateway"

```

GatewayToInternet:

```

Type: "AWS::EC2::VPCEGatewayAttachment"
Properties:
  VpcId: !Ref VPC
  InternetGatewayId: !Ref InternetGateway

```

```
PublicRouteTable:
  Type: "AWS::EC2::RouteTable"
  Properties:
    VpId: !Ref VPC
PublicRoute:
  Type: "AWS::EC2::Route"
  DependsOn: GatewayToInternet
  Properties:
    RouteTableId: !Ref PublicRouteTable
    DestinationCidrBlock: 0.0.0.0/0
    GatewayId: !Ref InternetGateway
PublicSubnetRouteTableAssociation:
  Type: "AWS::EC2::SubnetRouteTableAssociation"
  Properties:
    SubnetId: !Ref PublicSubnet
    RouteTableId: !Ref PublicRouteTable
PublicSubnetRouteTableAssociation2:
  Type: "AWS::EC2::SubnetRouteTableAssociation"
  Condition: DoAz2
  Properties:
    SubnetId: !Ref PublicSubnet2
    RouteTableId: !Ref PublicRouteTable
PublicSubnetRouteTableAssociation3:
  Condition: DoAz3
  Type: "AWS::EC2::SubnetRouteTableAssociation"
  Properties:
    SubnetId: !Ref PublicSubnet3
    RouteTableId: !Ref PublicRouteTable
PrivateSubnet:
  Type: "AWS::EC2::Subnet"
  Properties:
    VpId: !Ref VPC
    CidrBlock: !Select [3, !Cidr [!Ref VpcCidr, 6, !Ref SubnetBits]]
    AvailabilityZone: !Select
      - 0
      - Fn::GetAZs: !Ref "AWS::Region"
PrivateRouteTable:
  Type: "AWS::EC2::RouteTable"
  Properties:
    VpId: !Ref VPC
PrivateSubnetRouteTableAssociation:
  Type: "AWS::EC2::SubnetRouteTableAssociation"
  Properties:
    SubnetId: !Ref PrivateSubnet
    RouteTableId: !Ref PrivateRouteTable
NAT:
  DependsOn:
    - GatewayToInternet
  Type: "AWS::EC2::NatGateway"
  Properties:
    AllocationId:
      "Fn::GetAtt":
        - EIP
        - AllocationId
    SubnetId: !Ref PublicSubnet
EIP:
```



```

Type: "AWS::EC2::EIP"
Properties:
  Domain: vpc
Route:
  Type: "AWS::EC2::Route"
  Properties:
    RouteTableId:
      Ref: PrivateRouteTable
    DestinationCidrBlock: 0.0.0.0/0
    NatGatewayId:
      Ref: NAT
PrivateSubnet2:
  Type: "AWS::EC2::Subnet"
  Condition: DoAz2
  Properties:
    VpcId: !Ref VPC
    CidrBlock: !Select [4, !Cidr [!Ref VpcCidr, 6, !Ref SubnetBits]]
    AvailabilityZone: !Select
      - 1
      - Fn::GetAZs: !Ref "AWS::Region"
PrivateRouteTable2:
  Type: "AWS::EC2::RouteTable"
  Condition: DoAz2
  Properties:
    VpcId: !Ref VPC
PrivateSubnetRouteTableAssociation2:
  Type: "AWS::EC2::SubnetRouteTableAssociation"
  Condition: DoAz2
  Properties:
    SubnetId: !Ref PrivateSubnet2
    RouteTableId: !Ref PrivateRouteTable2
NAT2:
  DependsOn:
    - GatewayToInternet
  Type: "AWS::EC2::NatGateway"
  Condition: DoAz2
  Properties:
    AllocationId:
      "Fn::GetAtt":
        - EIP2
        - AllocationId
    SubnetId: !Ref PublicSubnet2
EIP2:
  Type: "AWS::EC2::EIP"
  Condition: DoAz2
  Properties:
    Domain: vpc
Route2:
  Type: "AWS::EC2::Route"
  Condition: DoAz2
  Properties:
    RouteTableId:
      Ref: PrivateRouteTable2
    DestinationCidrBlock: 0.0.0.0/0
    NatGatewayId:
      Ref: NAT2

```

```
PrivateSubnet3:
  Type: "AWS::EC2::Subnet"
  Condition: DoAz3
  Properties:
    VpcId: !Ref VPC
    CidrBlock: !Select [5, !Cidr [!Ref VpcCidr, 6, !Ref SubnetBits]]
    AvailabilityZone: !Select
      - 2
      - Fn::GetAZs: !Ref "AWS::Region"
PrivateRouteTable3:
  Type: "AWS::EC2::RouteTable"
  Condition: DoAz3
  Properties:
    VpcId: !Ref VPC
PrivateSubnetRouteTableAssociation3:
  Type: "AWS::EC2::SubnetRouteTableAssociation"
  Condition: DoAz3
  Properties:
    SubnetId: !Ref PrivateSubnet3
    RouteTableId: !Ref PrivateRouteTable3
NAT3:
  DependsOn:
    - GatewayToInternet
  Type: "AWS::EC2::NatGateway"
  Condition: DoAz3
  Properties:
    AllocationId:
      "Fn::GetAtt":
        - EIP3
        - AllocationId
    SubnetId: !Ref PublicSubnet3
EIP3:
  Type: "AWS::EC2::EIP"
  Condition: DoAz3
  Properties:
    Domain: vpc
Route3:
  Type: "AWS::EC2::Route"
  Condition: DoAz3
  Properties:
    RouteTableId:
      Ref: PrivateRouteTable3
    DestinationCidrBlock: 0.0.0.0/0
    NatGatewayId:
      Ref: NAT3
S3Endpoint:
  Type: AWS::EC2::VPCEndpoint
  Properties:
    PolicyDocument:
      Version: 2012-10-17
      Statement:
        - Effect: Allow
          Principal: '*'
          Action:
            - '*'
          Resource:
```

```

- !*!
RouteTableIds:
- !Ref PublicRouteTable
- !Ref PrivateRouteTable
- !If [DoAz2, !Ref PrivateRouteTable2, !Ref "AWS::NoValue"]
- !If [DoAz3, !Ref PrivateRouteTable3, !Ref "AWS::NoValue"]
ServiceName: !Join
- "
- - com.amazonaws.
- !Ref 'AWS::Region'
- .s3
VpId: !Ref VPC

Outputs:
VpId:
Description: ID of the new VPC.
Value: !Ref VPC
PublicSubnetIds:
Description: Subnet IDs of the public subnets.
Value:
!Join [
  ",",
  [!Ref PublicSubnet, !If [DoAz2, !Ref PublicSubnet2, !Ref "AWS::NoValue"], !If [DoAz3, !Ref
PublicSubnet3, !Ref "AWS::NoValue"]]
]
PrivateSubnetIds:
Description: Subnet IDs of the private subnets.
Value:
!Join [
  ",",
  [!Ref PrivateSubnet, !If [DoAz2, !Ref PrivateSubnet2, !Ref "AWS::NoValue"], !If [DoAz3, !Ref
PrivateSubnet3, !Ref "AWS::NoValue"]]
]
PublicRouteTableId:
Description: Public Route table ID
Value: !Ref PublicRouteTable
PrivateRouteTableIds:
Description: Private Route table IDs
Value:
!Join [
  ",",
  [
    !Join ["=", [
      !Select [0, "Fn::GetAZs": !Ref "AWS::Region"],
      !Ref PrivateRouteTable
    ]],
    !If [DoAz2,
      !Join ["=", [!Select [1, "Fn::GetAZs": !Ref "AWS::Region"], !Ref PrivateRouteTable2]],
      !Ref "AWS::NoValue"
    ],
    !If [DoAz3,
      !Join ["=", [!Select [2, "Fn::GetAZs": !Ref "AWS::Region"], !Ref PrivateRouteTable3]],
      !Ref "AWS::NoValue"
    ]
  ]
]
]
]

```

6.3.12.7.3. Creating a VPC carrier gateway

To use public subnets in your OpenShift Container Platform cluster that runs on Wavelength Zones, you must create the carrier gateway and associate the carrier gateway to the VPC. Subnets are useful for deploying load balancers or edge compute nodes.

To create edge nodes or internet-facing load balancers in Wavelength Zones locations for your OpenShift Container Platform cluster, you must create the following required network components:

- A carrier gateway that associates to the existing VPC.
- A carrier route table that lists route entries.
- A subnet that associates to the carrier route table.

Carrier gateways exist for VPCs that only contain subnets in a Wavelength Zone.

The following list explains the functions of a carrier gateway in the context of an AWS Wavelength Zones location:

- Provides connectivity between your Wavelength Zone and the carrier network, which includes any available devices from the carrier network.
- Performs Network Address Translation (NAT) functions, such as translating IP addresses that are public IP addresses stored in a network border group, from Wavelength Zones to carrier IP addresses. These translation functions apply to inbound and outbound traffic.
- Authorizes inbound traffic from a carrier network that is located in a specific location.
- Authorizes outbound traffic to a carrier network and the internet.



NOTE

No inbound connection configuration exists from the internet to a Wavelength Zone through the carrier gateway.

You can use the provided CloudFormation template to create a stack of the following AWS resources:

- One carrier gateway that associates to the VPC ID in the template.
- One public route table for the Wavelength Zone named as **<ClusterName>-public-carrier**.
- Default IPv4 route entry in the new route table that targets the carrier gateway.
- VPC gateway endpoint for an AWS Simple Storage Service (S3).



NOTE

If you do not use the provided CloudFormation template to create your AWS infrastructure, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- You configured an AWS account.
- You added your AWS keys and region to your local AWS profile by running **aws configure**.

Procedure

1. Go to the next section of the documentation named "CloudFormation template for the VPC Carrier Gateway", and then copy the syntax from the **CloudFormation template for VPC Carrier Gateway** template. Save the copied template syntax as a YAML file on your local system. This template describes the VPC that your cluster requires.
2. Run the following command to deploy the CloudFormation template, which creates a stack of AWS resources that represent the VPC:

```
$ aws cloudformation create-stack --stack-name <stack_name> \ 1
--region ${CLUSTER_REGION} \
--template-body file://<template>.yaml \ 2
--parameters \
ParameterKey=VpcId,ParameterValue="${VpcId}" \ 3
ParameterKey=ClusterName,ParameterValue="${ClusterName}" \ 4
```

- 1 **<stack_name>** is the name for the CloudFormation stack, such as **clusterName-vpc-carrier-gw**. You need the name of this stack if you remove the cluster.
- 2 **<template>** is the relative path and the name of the CloudFormation template YAML file that you saved.
- 3 **<VpcId>** is the VPC ID extracted from the CloudFormation stack output created in the section named "Creating a VPC in AWS".
- 4 **<ClusterName>** is a custom value that prefixes to resources that the CloudFormation stack creates. You can use the same name that is defined in the **metadata.name** section of the **install-config.yaml** configuration file.

Example output

```
arn:aws:cloudformation:us-east-1:123456789012:stack/<stack_name>/dbedae40-2fd3-11eb-820e-12a48460849f
```

Verification

- Confirm that the CloudFormation template components exist by running the following command:

```
$ aws cloudformation describe-stacks --stack-name <stack_name>
```

After the **StackStatus** displays **CREATE_COMPLETE**, the output displays values for the following parameter. Ensure that you provide the parameter value to the other CloudFormation templates that you run to create for your cluster.

PublicRouteTableId

The ID of the Route Table in the Carrier infrastructure.

Additional resources

- See [Amazon S3](#) in the AWS documentation.

6.3.12.7.4. CloudFormation template for the VPC Carrier Gateway

You can use the following CloudFormation template to deploy the Carrier Gateway on AWS Wavelength infrastructure.

Example 6.53. CloudFormation template for VPC Carrier Gateway

```

AWSTemplateFormatVersion: 2010-09-09
Description: Template for Creating Wavelength Zone Gateway (Carrier Gateway).

Parameters:
  VpcId:
    Description: VPC ID to associate the Carrier Gateway.
    Type: String
    AllowedPattern: ^(?::(?:vpc)(?:-[a-zA-Z0-9]+)?\b|(?:[0-9]{1,3}\.){3}[0-9]{1,3})$
    ConstraintDescription: VPC ID must be with valid name, starting with vpc-.*
  ClusterName:
    Description: Cluster Name or Prefix name to prepend the tag Name for each subnet.
    Type: String
    AllowedPattern: ".+"
    ConstraintDescription: ClusterName parameter must be specified.

Resources:
  CarrierGateway:
    Type: "AWS::EC2::CarrierGateway"
    Properties:
      VpcId: !Ref VpcId
      Tags:
        - Key: Name
          Value: !Join ['-', [!Ref ClusterName, "cagw"]]

  PublicRouteTable:
    Type: "AWS::EC2::RouteTable"
    Properties:
      VpcId: !Ref VpcId
      Tags:
        - Key: Name
          Value: !Join ['-', [!Ref ClusterName, "public-carrier"]]

  PublicRoute:
    Type: "AWS::EC2::Route"
    DependsOn: CarrierGateway
    Properties:
      RouteTableId: !Ref PublicRouteTable
      DestinationCidrBlock: 0.0.0.0/0
      CarrierGatewayId: !Ref CarrierGateway

```

```

S3Endpoint:
  Type: AWS::EC2::VPCEndpoint
  Properties:
    PolicyDocument:
      Version: 2012-10-17
      Statement:
        - Effect: Allow
          Principal: '*'
          Action:
            - '*'
          Resource:
            - '*'

    RouteTableIds:
      - !Ref PublicRouteTable
    ServiceName: !Join
      - ""
      - - com.amazonaws.
        - !Ref 'AWS::Region'
        - .s3
    VpId: !Ref VpId

```

```


Outputs:
  PublicRouteTableId:
    Description: Public Route table ID
    Value: !Ref PublicRouteTable

```

6.3.12.7.5. Creating subnets in Wavelength Zones

Before you configure a machine set for edge compute nodes in your OpenShift Container Platform cluster, you must create the subnets in Wavelength Zones. Complete the following procedure for each Wavelength Zone that you want to deploy compute nodes to.

You can use the provided CloudFormation template and create a CloudFormation stack. You can then use this stack to custom provision a subnet.



NOTE

If you do not use the provided CloudFormation template to create your AWS infrastructure, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- You configured an AWS account.
- You added your AWS keys and region to your local AWS profile by running **aws configure**.
- You opted in to the Wavelength Zones group.

Procedure

1. Go to the section of the documentation named "CloudFormation template for the VPC subnet", and copy the syntax from the template. Save the copied template syntax as a YAML file on your local system. This template describes the VPC that your cluster requires.
2. Run the following command to deploy the CloudFormation template, which creates a stack of AWS resources that represent the VPC:

```
$ aws cloudformation create-stack --stack-name <stack_name> \ 1
--region ${CLUSTER_REGION} \
--template-body file://<template>.yaml \ 2
--parameters \
ParameterKey=VpcId,ParameterValue="${VPC_ID}" \ 3
ParameterKey=ClusterName,ParameterValue="${CLUSTER_NAME}" \ 4
ParameterKey=ZoneName,ParameterValue="${ZONE_NAME}" \ 5
ParameterKey=PublicRouteTableId,ParameterValue="${ROUTE_TABLE_PUB}" \ 6
ParameterKey=PublicSubnetCidr,ParameterValue="${SUBNET_CIDR_PUB}" \ 7
ParameterKey=PrivateRouteTableId,ParameterValue="${ROUTE_TABLE_PVT}" \ 8
ParameterKey=PrivateSubnetCidr,ParameterValue="${SUBNET_CIDR_PVT}" \ 9
```

- 1 **<stack_name>** is the name for the CloudFormation stack, such as **cluster-wl-
<wavelength_zone_shortcode>**. You need the name of this stack if you remove the cluster.
- 2 **<template>** is the relative path and the name of the CloudFormation template YAML file that you saved.
- 3 **\${VPC_ID}** is the VPC ID, which is the value **VpcId** in the output of the CloudFormation template for the VPC.
- 4 **\${ZONE_NAME}** is the value of Wavelength Zones name to create the subnets.
- 5 **\${CLUSTER_NAME}** is the value of **ClusterName** to be used as a prefix of the new AWS resource names.
- 6 **\${ROUTE_TABLE_PUB}** is the **PublicRouteTableId** extracted from the output of the VPC's carrier gateway CloudFormation stack.
- 7 **\${SUBNET_CIDR_PUB}** is a valid CIDR block that is used to create the public subnet. This block must be part of the VPC CIDR block **VpcCidr**.
- 8 **\${ROUTE_TABLE_PVT}** is the **PrivateRouteTableId** extracted from the output of the VPC's CloudFormation stack.
- 9 **\${SUBNET_CIDR_PVT}** is a valid CIDR block that is used to create the private subnet. This block must be part of the VPC CIDR block **VpcCidr**.

Example output

```
arn:aws:cloudformation:us-east-1:123456789012:stack/<stack_name>/dbedae40-820e-11eb-2fd3-12a48460849f
```

Verification

- Confirm that the template components exist by running the following command:

```
$ aws cloudformation describe-stacks --stack-name <stack_name>
```

After the **StackStatus** displays **CREATE_COMPLETE**, the output displays values for the following parameters. Ensure that you provide these parameter values to the other CloudFormation templates that you run to create for your cluster.

PublicSubnetId	The IDs of the public subnet created by the CloudFormation stack.
PrivateSubnetId	The IDs of the private subnet created by the CloudFormation stack.

6.3.12.7.6. CloudFormation template for the VPC subnet

You can use the following CloudFormation template to deploy the private and public subnets in a zone on Wavelength Zones infrastructure.

Example 6.54. CloudFormation template for VPC subnets

AWSTemplateFormatVersion: 2010-09-09

Description: Template for Best Practice Subnets (Public and Private)

Parameters:

VpcId:

Description: VPC ID that comprises all the target subnets.

Type: String

AllowedPattern: `^(?:(:vpc)(?:-[a-zA-Z0-9]+)?\b(?:[0-9]{1,3}\.){3}[0-9]{1,3})$`

ConstraintDescription: VPC ID must be with valid name, starting with vpc-.*.

ClusterName:

Description: Cluster name or prefix name to prepend the Name tag for each subnet.

Type: String

AllowedPattern: `".+"`

ConstraintDescription: ClusterName parameter must be specified.

ZoneName:

Description: Zone Name to create the subnets, such as us-west-2-lax-1a.

Type: String

AllowedPattern: `".+"`

ConstraintDescription: ZoneName parameter must be specified.

PublicRouteTableId:

Description: Public Route Table ID to associate the public subnet.

Type: String

AllowedPattern: `".+"`

ConstraintDescription: PublicRouteTableId parameter must be specified.

PublicSubnetCidr:

AllowedPattern: `^(((0-9){1,3}|1[0-9]{2}|2[0-4][0-9]|25[0-5])\.)\.(0-9){1,3}|1[0-9]{2}|2[0-4][0-9]|25[0-5])(\.(1[6-9]|2[0-4]))$`

ConstraintDescription: CIDR block parameter must be in the form x.x.x.x/16-24.

Default: `10.0.128.0/20`

Description: CIDR block for public subnet.

Type: String

PrivateRouteTableId:

Description: Private Route Table ID to associate the private subnet.

Type: String

AllowedPattern: ".+"

ConstraintDescription: PrivateRouteTableId parameter must be specified.

PrivateSubnetCidr:

AllowedPattern: ^(([0-9]|[1-9][0-9]|1[0-9]{2}|2[0-4][0-9]|25[0-5])\.)\{3\}([0-9]|[1-9][0-9]|1[0-9]{2}|2[0-4][0-9]|25[0-5])\(\{1[6-9]|2[0-4]\}\)\$

ConstraintDescription: CIDR block parameter must be in the form x.x.x.x/16-24.

Default: 10.0.128.0/20

Description: CIDR block for private subnet.

Type: String

Resources:

PublicSubnet:

Type: "AWS::EC2::Subnet"

Properties:

Vpclid: !Ref Vpclid

CidrBlock: !Ref PublicSubnetCidr

AvailabilityZone: !Ref ZoneName

Tags:

- Key: Name

Value: !Join ['-', [!Ref ClusterName, "public", !Ref ZoneName]]

PublicSubnetRouteTableAssociation:

Type: "AWS::EC2::SubnetRouteTableAssociation"

Properties:

SubnetId: !Ref PublicSubnet

RouteTableId: !Ref PublicRouteTableId

PrivateSubnet:

Type: "AWS::EC2::Subnet"

Properties:

Vpclid: !Ref Vpclid

CidrBlock: !Ref PrivateSubnetCidr

AvailabilityZone: !Ref ZoneName

Tags:

- Key: Name

Value: !Join ['-', [!Ref ClusterName, "private", !Ref ZoneName]]

PrivateSubnetRouteTableAssociation:

Type: "AWS::EC2::SubnetRouteTableAssociation"

Properties:

SubnetId: !Ref PrivateSubnet

RouteTableId: !Ref PrivateRouteTableId

Outputs:

PublicSubnetId:

Description: Subnet ID of the public subnets.

Value:

!Join [",", [!Ref PublicSubnet]]

PrivateSubnetId:

Description: Subnet ID of the private subnets.

Value:

!Join [",", [!Ref PrivateSubnet]]

6.3.12.7.7. Modifying an installation configuration file to use AWS Wavelength Zones subnets

Modify your **install-config.yaml** file to include Wavelength Zones subnets.

Prerequisites

- You created subnets by using the procedure "Creating subnets in Wavelength Zones".
- You created an **install-config.yaml** file by using the procedure "Creating the installation configuration file".

Procedure

- Modify the **install-config.yaml** configuration file by specifying Wavelength Zones subnets in the **platform.aws.subnets** parameter.

Example installation configuration file with Wavelength Zones subnets

```
# ...
platform:
  aws:
    region: us-west-2
    subnets: ❶
    - publicSubnetId-1
    - publicSubnetId-2
    - publicSubnetId-3
    - privateSubnetId-1
    - privateSubnetId-2
    - privateSubnetId-3
    - publicOrPrivateSubnetID-Wavelength-1
# ...
```

- ❶ List of subnet IDs created in the zones: Availability and Wavelength Zones.

Additional resources

- For more information about viewing the CloudFormation stacks that you created, see [AWS CloudFormation console](#).
- For more information about AWS profile and credential configuration, see [Configuration and credential file settings](#) in the AWS documentation.

Next steps

- [Deploying the cluster](#)

6.3.12.8. Optional: Assign public IP addresses to edge compute nodes

If your workload requires deploying the edge compute nodes in public subnets on Wavelength Zones infrastructure, you can configure the machine set manifests when installing a cluster.

AWS Wavelength Zones infrastructure accesses the network traffic in a specified zone, so applications can take advantage of lower latency when serving end users that are closer to that zone.

The default setting that deploys compute nodes in private subnets might not meet your needs, so consider creating edge compute nodes in public subnets when you want to apply more customization to your infrastructure.



IMPORTANT

By default, OpenShift Container Platform deploys the compute nodes in private subnets. For best performance, consider placing compute nodes in subnets that have their Public IP addresses attached to the subnets.

You must create additional security groups, but ensure that you only open the groups' rules over the internet when you really need to.

Procedure

1. Change to the directory that contains the installation program and generate the manifest files. Ensure that the installation manifests get created at the **openshift** and **manifests** directory level.

```
$ ./openshift-install create manifests --dir <installation_directory>
```

2. Edit the machine set manifest that the installation program generates for the Wavelength Zones, so that the manifest gets deployed in public subnets. Specify **true** for the **spec.template.spec.providerSpec.value.publicIP** parameter.

Example machine set manifest configuration for installing a cluster quickly in Wavelength Zones

```
spec:
  template:
    spec:
      providerSpec:
        value:
          publicip: true
          subnet:
            filters:
              - name: tag:Name
            values:
              - ${INFRA_ID}-public-${ZONE_NAME}
```

Example machine set manifest configuration for installing a cluster in an existing VPC that has Wavelength Zones subnets

```
apiVersion: machine.openshift.io/v1beta1
kind: MachineSet
metadata:
  name: <infrastructure_id>-edge-<zone>
  namespace: openshift-machine-api
spec:
  template:
    spec:
```

```
providerSpec:
  value:
    publicIp: true
```

6.3.12.9. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.

Procedure

1. Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

- 1** For **<installation_directory>**, specify the location of your customized **./install-config.yaml** file.
- 2** To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

2. Optional: Remove or disable the **AdministratorAccess** policy from the IAM account that you used to install the cluster.



NOTE

The elevated permissions provided by the **AdministratorAccess** policy are required only during installation.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to **<installation_directory>/openshift_install.log**.

**IMPORTANT**

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```

**IMPORTANT**

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

6.3.12.10. Verifying the status of the deployed cluster

Verify that your OpenShift Container Platform successfully deployed on AWS Wavelength Zones.

6.3.12.10.1. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

6.3.12.10.2. Logging in to the cluster by using the web console

The **kubeadmin** user exists by default after an OpenShift Container Platform installation. You can log in to your cluster as the **kubeadmin** user by using the OpenShift Container Platform web console.

Prerequisites

- You have access to the installation host.
- You completed a cluster installation and all cluster Operators are available.

Procedure

1. Obtain the password for the **kubeadmin** user from the **kubeadmin-password** file on the installation host:

```
$ cat <installation_directory>/auth/kubeadmin-password
```




NOTE

Alternatively, you can obtain the **kubeadmin** password from the **<installation_directory>/openshift_install.log** log file on the installation host.

2. List the OpenShift Container Platform web console route:

```
$ oc get routes -n openshift-console | grep 'console-openshift'
```



NOTE

Alternatively, you can obtain the OpenShift Container Platform route from the **<installation_directory>/openshift_install.log** log file on the installation host.

Example output

```
console    console-openshift-console.apps.<cluster_name>.<base_domain>    console
https reencrypt/Redirect None
```

3. Navigate to the route detailed in the output of the preceding command in a web browser and log in as the **kubeadmin** user.

Additional resources

- [Accessing the web console](#)

6.3.12.10.3. Verifying nodes that were created with edge compute pool

After you install a cluster that uses AWS Wavelength Zones infrastructure, check the status of the machine that was created by the machine set manifests created during installation.

1. To check the machine sets created from the subnet you added to the **install-config.yaml** file, run the following command:

```
$ oc get machineset -n openshift-machine-api
```

Example output

NAME	DESIRED	CURRENT	READY	AVAILABLE	AGE
cluster-7xw5g-edge-us-east-1-wl1-nyc-wlz-1	1	1	1	1	3h4m
cluster-7xw5g-worker-us-east-1a	1	1	1	1	3h4m
cluster-7xw5g-worker-us-east-1b	1	1	1	1	3h4m
cluster-7xw5g-worker-us-east-1c	1	1	1	1	3h4m

2. To check the machines that were created from the machine sets, run the following command:

```
$ oc get machines -n openshift-machine-api
```

Example output

NAME	PHASE	TYPE	REGION	ZONE	AGE
cluster-7xw5g-edge-us-east-1-wl1-nyc-wlz-1-wbclh	Running		c5d.2xlarge	us-east-1	us-east-1-wl1-nyc-wlz-1 3h
cluster-7xw5g-master-0	Running	m6i.xlarge	us-east-1	us-east-1a	3h4m
cluster-7xw5g-master-1	Running	m6i.xlarge	us-east-1	us-east-1b	3h4m
cluster-7xw5g-master-2	Running	m6i.xlarge	us-east-1	us-east-1c	3h4m
cluster-7xw5g-worker-us-east-1a-rtp45	Running	m6i.xlarge	us-east-1	us-east-1a	3h
cluster-7xw5g-worker-us-east-1b-glm7c	Running	m6i.xlarge	us-east-1	us-east-1b	3h
cluster-7xw5g-worker-us-east-1c-qfvz4	Running	m6i.xlarge	us-east-1	us-east-1c	3h

3. To check nodes with edge roles, run the following command:

```
$ oc get nodes -l node-role.kubernetes.io/edge
```

Example output

NAME	STATUS	ROLES	AGE	VERSION
ip-10-0-207-188.ec2.internal	Ready	edge,worker	172m	v1.25.2+d2e245f

Next steps

- [Validating an installation.](#)
- If necessary, you can [opt out of remote health](#).

6.3.13. Installing a cluster with compute nodes on AWS Outposts

In OpenShift Container Platform version 4.14, you could install a cluster on Amazon Web Services (AWS) with compute nodes running in AWS Outposts as a Technology Preview. As of OpenShift Container Platform version 4.15, this installation method is no longer supported.

Instead, you can [install a cluster on AWS into an existing VPC](#) and provision compute nodes on AWS Outposts as a postinstallation configuration task.

For more information, see [Extending an AWS VPC cluster into an AWS Outpost](#)

6.4. USER-PROVISIONED INFRASTRUCTURE

6.4.1. Preparing to install a cluster on AWS

You prepare to install an OpenShift Container Platform cluster on AWS by completing the following steps:

- Verifying internet connectivity for your cluster.
- [Configuring an AWS account.](#)
- Downloading the installation program.



NOTE

If you are installing in a disconnected environment, you extract the installation program from the mirrored content. For more information, see [Mirroring images for a disconnected installation](#).

- Installing the OpenShift CLI (**oc**).



NOTE

If you are installing in a disconnected environment, install **oc** to the mirror host.

- Generating an SSH key pair. You can use this key pair to authenticate into the OpenShift Container Platform cluster's nodes after it is deployed.
- [Preparing the user-provisioned infrastructure.](#)
- If the cloud identity and access management (IAM) APIs are not accessible in your environment, or if you do not want to store an administrator-level credential secret in the **kube-system** namespace, [manually creating long-term credentials for AWS](#) or [configuring an AWS cluster to](#)

use [short-term credentials](#) with Amazon Web Services Security Token Service (AWS STS).

6.4.1.1. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.



IMPORTANT

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

6.4.1.2. Obtaining the installation program

Before you install OpenShift Container Platform, download the installation file on the host you are using for installation.

Prerequisites

- You have a computer that runs Linux or macOS, with 500 MB of local disk space.


Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select your infrastructure provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.



IMPORTANT

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.



IMPORTANT

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.


4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#). This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

6.4.1.3. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.



IMPORTANT

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**. To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.



NOTE

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your **PATH**.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

6.4.1.4. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.



NOTE

You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.



NOTE

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.


2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.



NOTE

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the `ssh-agent` process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```



NOTE

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the `ssh-agent`:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

6.4.1.5. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service.

6.4.2. Installation requirements for user-provisioned infrastructure on AWS

Before you begin an installation on infrastructure that you provision, be sure that your AWS environment meets the following installation requirements.

For a cluster that contains user-provisioned infrastructure, you must deploy all of the required machines.

6.4.2.1. Required machines for cluster installation

The smallest OpenShift Container Platform clusters require the following hosts:

Table 6.21. Minimum required hosts

Hosts	Description
One temporary bootstrap machine	The cluster requires the bootstrap machine to deploy the OpenShift Container Platform cluster on the three control plane machines. You can remove the bootstrap machine after you install the cluster.
Three control plane machines	The control plane machines run the Kubernetes and OpenShift Container Platform services that form the control plane.
At least two compute machines, which are also known as worker machines.	The workloads requested by OpenShift Container Platform users run on the compute machines.

**IMPORTANT**

To maintain high availability of your cluster, use separate physical hosts for these cluster machines.

The bootstrap and control plane machines must use Red Hat Enterprise Linux CoreOS (RHCOS) as the operating system. However, the compute machines can choose between Red Hat Enterprise Linux CoreOS (RHCOS), Red Hat Enterprise Linux (RHEL) 8.6 and later.

Note that RHCOS is based on Red Hat Enterprise Linux (RHEL) 9.2 and inherits all of its hardware certifications and requirements. See [Red Hat Enterprise Linux technology capabilities and limits](#) .

6.4.2.1.1. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 6.22. Minimum resource requirements

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Bootstrap	RHCOS	4	16 GB	100 GB	300
Control plane	RHCOS	4	16 GB	100 GB	300
Compute	RHCOS, RHEL 8.6 and later [3]	2	8 GB	100 GB	300

1. One vCPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio: $(\text{threads per core} \times \text{cores}) \times \text{sockets} = \text{vCPUs}$.
2. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes which require a 10 ms p99 fsync duration. Note that on many cloud platforms, storage size and IOPS scale together, so you might need to over-allocate storage volume to obtain sufficient performance.
3. As with all user-provisioned installations, if you choose to use RHEL compute machines in your cluster, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Use of RHEL 7 compute machines is deprecated and has been removed in OpenShift Container Platform 4.10 and later.

NOTE

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).


If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

6.4.2.1.2. Tested instance types for AWS

The following Amazon Web Services (AWS) instance types have been tested with OpenShift Container Platform.



NOTE

Use the machine types included in the following charts for your AWS instances. If you use an instance type that is not listed in the chart, ensure that the instance size you use matches the minimum resource requirements that are listed in the section named "Minimum resource requirements for cluster installation".

Example 6.55. Machine types based on 64-bit x86 architecture

- **c4.***
- **c5.***
- **c5a.***
- **i3.***
- **m4.***
- **m5.***
- **m5a.***
- **m6a.***
- **m6i.***
- **r4.***
- **r5.***
- **r5a.***
- **r6i.***
- **t3.***
- **t3a.***

6.4.2.1.3. Tested instance types for AWS on 64-bit ARM infrastructures

The following Amazon Web Services (AWS) 64-bit ARM instance types have been tested with OpenShift Container Platform.



NOTE

Use the machine types included in the following charts for your AWS ARM instances. If you use an instance type that is not listed in the chart, ensure that the instance size you use matches the minimum resource requirements that are listed in "Minimum resource requirements for cluster installation".

Example 6.56. Machine types based on 64-bit ARM architecture

- **c6g.***
- **m6g.***

6.4.2.2. Certificate signing requests management

Because your cluster has limited access to automatic machine management when you use infrastructure that you provision, you must provide a mechanism for approving cluster certificate signing requests (CSRs) after installation. The **kube-controller-manager** only approves the kubelet client CSRs. The **machine-approver** cannot guarantee the validity of a serving certificate that is requested by using kubelet credentials because it cannot confirm that the correct machine issued the request. You must determine and implement a method of verifying the validity of the kubelet serving certificate requests and approving them.

6.4.2.3. Required AWS infrastructure components

To install OpenShift Container Platform on user-provisioned infrastructure in Amazon Web Services (AWS), you must manually create both the machines and their supporting infrastructure.

For more information about the integration testing for different platforms, see the [OpenShift Container Platform 4.x Tested Integrations](#) page.

By using the provided CloudFormation templates, you can create stacks of AWS resources that represent the following components:

- An AWS Virtual Private Cloud (VPC)
- Networking and load balancing components
- Security groups and roles
- An OpenShift Container Platform bootstrap node
- OpenShift Container Platform control plane nodes
- An OpenShift Container Platform compute node

Alternatively, you can manually create the components or you can reuse existing infrastructure that meets the cluster requirements. Review the CloudFormation templates for more details about how the components interrelate.

6.4.2.3.1. Other infrastructure components

- A VPC
- DNS entries
- Load balancers (classic or network) and listeners
- A public and a private Route 53 zone
- Security groups

- IAM roles
- S3 buckets

If you are working in a disconnected environment, you are unable to reach the public IP addresses for EC2, ELB, and S3 endpoints. Depending on the level to which you want to restrict internet traffic during the installation, the following configuration options are available:

Option 1: Create VPC endpoints

Create a VPC endpoint and attach it to the subnets that the clusters are using. Name the endpoints as follows:

- **ec2.<aws_region>.amazonaws.com**
- **elasticloadbalancing.<aws_region>.amazonaws.com**
- **s3.<aws_region>.amazonaws.com**

With this option, network traffic remains private between your VPC and the required AWS services.

Option 2: Create a proxy without VPC endpoints

As part of the installation process, you can configure an HTTP or HTTPS proxy. With this option, internet traffic goes through the proxy to reach the required AWS services.

Option 3: Create a proxy with VPC endpoints

As part of the installation process, you can configure an HTTP or HTTPS proxy with VPC endpoints. Create a VPC endpoint and attach it to the subnets that the clusters are using. Name the endpoints as follows:

- **ec2.<aws_region>.amazonaws.com**
- **elasticloadbalancing.<aws_region>.amazonaws.com**
- **s3.<aws_region>.amazonaws.com**

When configuring the proxy in the **install-config.yaml** file, add these endpoints to the **noProxy** field. With this option, the proxy prevents the cluster from accessing the internet directly. However, network traffic remains private between your VPC and the required AWS services.

Required VPC components

You must provide a suitable VPC and subnets that allow communication to your machines.

Component	AWS type	Description
VPC	<ul style="list-style-type: none"> • AWS::EC2::VPC • AWS::EC2::VPCEndpoint 	You must provide a public VPC for the cluster to use. The VPC uses an endpoint that references the route tables for each subnet to improve communication with the registry that is hosted in S3.

Component	AWS type	Description												
Public subnets	<ul style="list-style-type: none"> ● AWS::EC2::Subnet ● AWS::EC2::SubnetNetworkACLAssociation 	Your VPC must have public subnets for between 1 and 3 availability zones and associate them with appropriate Ingress rules.												
Internet gateway	<ul style="list-style-type: none"> ● AWS::EC2::InternetGateway ● AWS::EC2::VPCGatewayAttachment ● AWS::EC2::RouteTable ● AWS::EC2::Route ● AWS::EC2::SubnetRouteTableAssociation ● AWS::EC2::NatGateway ● AWS::EC2::EIP 	You must have a public internet gateway, with public routes, attached to the VPC. In the provided templates, each public subnet has a NAT gateway with an EIP address. These NAT gateways allow cluster resources, like private subnet instances, to reach the internet and are not required for some restricted network or proxy scenarios.												
Network access control	<ul style="list-style-type: none"> ● AWS::EC2::NetworkACL ● AWS::EC2::NetworkACLEntry 	<p>You must allow the VPC to access the following ports:</p> <table border="1"> <thead> <tr> <th>Port</th> <th>Reason</th> </tr> </thead> <tbody> <tr> <td>80</td> <td>Inbound HTTP traffic</td> </tr> <tr> <td>443</td> <td>Inbound HTTPS traffic</td> </tr> <tr> <td>22</td> <td>Inbound SSH traffic</td> </tr> <tr> <td>1024 - 65535</td> <td>Inbound ephemeral traffic</td> </tr> <tr> <td>0 - 65535</td> <td>Outbound ephemeral traffic</td> </tr> </tbody> </table>	Port	Reason	80	Inbound HTTP traffic	443	Inbound HTTPS traffic	22	Inbound SSH traffic	1024 - 65535	Inbound ephemeral traffic	0 - 65535	Outbound ephemeral traffic
Port	Reason													
80	Inbound HTTP traffic													
443	Inbound HTTPS traffic													
22	Inbound SSH traffic													
1024 - 65535	Inbound ephemeral traffic													
0 - 65535	Outbound ephemeral traffic													
Private subnets	<ul style="list-style-type: none"> ● AWS::EC2::Subnet ● AWS::EC2::RouteTable ● AWS::EC2::SubnetRouteTableAssociation 	Your VPC can have private subnets. The provided CloudFormation templates can create private subnets for between 1 and 3 availability zones. If you use private subnets, you must provide appropriate routes and tables for them.												

Required DNS and load balancing components

Your DNS and load balancer configuration needs to use a public hosted zone and can use a private hosted zone similar to the one that the installation program uses if it provisions the cluster's infrastructure. You must create a DNS entry that resolves to your load balancer. An entry for **api.<cluster_name>.<domain>** must point to the external load balancer, and an entry for **api-int.<cluster_name>.<domain>** must point to the internal load balancer.

The cluster also requires load balancers and listeners for port 6443, which are required for the Kubernetes API and its extensions, and port 22623, which are required for the Ignition config files for new machines. The targets will be the control plane nodes. Port 6443 must be accessible to both clients external to the cluster and nodes within the cluster. Port 22623 must be accessible to nodes within the cluster.

Component	AWS type	Description
DNS	AWS::Route53::HostedZone	The hosted zone for your internal DNS.
Public load balancer	AWS::ElasticLoadBalancingV2::LoadBalancer	The load balancer for your public subnets.
External API server record	AWS::Route53::RecordSetGroup	Alias records for the external API server.
External listener	AWS::ElasticLoadBalancingV2::Listener	A listener on port 6443 for the external load balancer.
External target group	AWS::ElasticLoadBalancingV2::TargetGroup	The target group for the external load balancer.
Private load balancer	AWS::ElasticLoadBalancingV2::LoadBalancer	The load balancer for your private subnets.
Internal API server record	AWS::Route53::RecordSetGroup	Alias records for the internal API server.
Internal listener	AWS::ElasticLoadBalancingV2::Listener	A listener on port 22623 for the internal load balancer.

Component	AWS type	Description
Internal target group	AWS::ElasticLoadBalancingV2::TargetGroup	The target group for the internal load balancer.
Internal listener	AWS::ElasticLoadBalancingV2::Listener	A listener on port 6443 for the internal load balancer.
Internal target group	AWS::ElasticLoadBalancingV2::TargetGroup	The target group for the internal load balancer.

Security groups

The control plane and worker machines require access to the following ports:

Group	Type	IP Protocol	Port range
MasterSecurityGroup	AWS::EC2::SecurityGroup	icmp	0
		tcp	22
		tcp	6443
		tcp	22623
WorkerSecurityGroup	AWS::EC2::SecurityGroup	icmp	0
		tcp	22
BootstrapSecurityGroup	AWS::EC2::SecurityGroup	tcp	22
		tcp	19531

Control plane Ingress

The control plane machines require the following Ingress groups. Each Ingress group is a **AWS::EC2::SecurityGroupIngress** resource.

Ingress group	Description	IP protocol	Port range
MasterIngressEtcd	etcd	tcp	2379- 2380

Ingress group	Description	IP protocol	Port range
MasterIngress Vxlan	Vxlan packets	udp	4789
MasterIngress WorkerVxlan	Vxlan packets	udp	4789
MasterIngress Internal	Internal cluster communication and Kubernetes proxy metrics	tcp	9000 - 9999
MasterIngress WorkerInternal	Internal cluster communication	tcp	9000 - 9999
MasterIngress Kube	Kubernetes kubelet, scheduler and controller manager	tcp	10250 - 10259
MasterIngress WorkerKube	Kubernetes kubelet, scheduler and controller manager	tcp	10250 - 10259
MasterIngress IngressServices	Kubernetes Ingress services	tcp	30000 - 32767
MasterIngress WorkerIngressServices	Kubernetes Ingress services	tcp	30000 - 32767
MasterIngress Geneve	Geneve packets	udp	6081
MasterIngress WorkerGeneve	Geneve packets	udp	6081
MasterIngress IpsecIke	IPsec IKE packets	udp	500
MasterIngress WorkerIpsecIke	IPsec IKE packets	udp	500
MasterIngress IpsecNat	IPsec NAT-T packets	udp	4500

Ingress group	Description	IP protocol	Port range
MasterIngress WorkerIpsecNat	IPsec NAT-T packets	udp	4500
MasterIngress IpsecEsp	IPsec ESP packets	50	All
MasterIngress WorkerIpsecEsp	IPsec ESP packets	50	All
MasterIngress InternalUDP	Internal cluster communication	udp	9000 - 9999
MasterIngress WorkerInternalUDP	Internal cluster communication	udp	9000 - 9999
MasterIngress IngressServicesUDP	Kubernetes Ingress services	udp	30000 - 32767
MasterIngress WorkerIngressServicesUDP	Kubernetes Ingress services	udp	30000 - 32767

Worker Ingress

The worker machines require the following Ingress groups. Each Ingress group is a **AWS::EC2::SecurityGroupIngress** resource.

Ingress group	Description	IP protocol	Port range
WorkerIngress Vxlan	Vxlan packets	udp	4789
WorkerIngress WorkerVxlan	Vxlan packets	udp	4789
WorkerIngress Internal	Internal cluster communication	tcp	9000 - 9999
WorkerIngress WorkerInternal	Internal cluster communication	tcp	9000 - 9999

Ingress group	Description	IP protocol	Port range
WorkerIngress Kube	Kubernetes kubelet, scheduler, and controller manager	tcp	10250
WorkerIngress WorkerKube	Kubernetes kubelet, scheduler, and controller manager	tcp	10250
WorkerIngress IngressServices	Kubernetes Ingress services	tcp	30000 - 32767
WorkerIngress WorkerIngress Services	Kubernetes Ingress services	tcp	30000 - 32767
WorkerIngress Geneve	Geneve packets	udp	6081
WorkerIngress MasterGeneve	Geneve packets	udp	6081
WorkerIngress IpsecIke	IPsec IKE packets	udp	500
WorkerIngress MasterIpsecIke	IPsec IKE packets	udp	500
WorkerIngress IpsecNat	IPsec NAT-T packets	udp	4500
WorkerIngress MasterIpsecNat	IPsec NAT-T packets	udp	4500
WorkerIngress IpsecEsp	IPsec ESP packets	50	All
WorkerIngress MasterIpsecEsp	IPsec ESP packets	50	All
WorkerIngress InternalUDP	Internal cluster communication	udp	9000 - 9999
WorkerIngress MasterInternalUDP	Internal cluster communication	udp	9000 - 9999

Ingress group	Description	IP protocol	Port range
WorkerIngressIngressServicesUDP	Kubernetes Ingress services	udp	30000 - 32767
WorkerIngressMasterIngressServicesUDP	Kubernetes Ingress services	udp	30000 - 32767

Roles and instance profiles

You must grant the machines permissions in AWS. The provided CloudFormation templates grant the machines **Allow** permissions for the following **AWS::IAM::Role** objects and provide a **AWS::IAM::InstanceProfile** for each set of roles. If you do not use the templates, you can grant the machines the following broad permissions or the following individual permissions.

Role	Effect	Action	Resource
Master	Allow	ec2:*	*
	Allow	elasticloadbalancing:*	*
	Allow	iam:PassRole	*
	Allow	s3:GetObject	*
Worker	Allow	ec2:Describe*	*
Bootstrap	Allow	ec2:Describe*	*
	Allow	ec2:AttachVolume	*
	Allow	ec2:DetachVolume	*

6.4.2.3.2. Cluster machines

You need **AWS::EC2::Instance** objects for the following machines:

- A bootstrap machine. This machine is required during installation, but you can remove it after your cluster deploys.
- Three control plane machines. The control plane machines are not governed by a control plane machine set.
- Compute machines. You must create at least two compute machines, which are also known as worker machines, during installation. These machines are not governed by a compute machine set.

6.4.2.4. Required AWS permissions for the IAM user



NOTE

Your IAM user must have the permission **tag:GetResources** in the region **us-east-1** to delete the base cluster resources. As part of the AWS API requirement, the OpenShift Container Platform installation program performs various actions in this region.

When you attach the **AdministratorAccess** policy to the IAM user that you create in Amazon Web Services (AWS), you grant that user all of the required permissions. To deploy all components of an OpenShift Container Platform cluster, the IAM user requires the following permissions:

Example 6.57. Required EC2 permissions for installation


- **ec2:AttachNetworkInterface**
- **ec2:AuthorizeSecurityGroupEgress**
- **ec2:AuthorizeSecurityGroupIngress**
- **ec2:CopyImage**
- **ec2>CreateNetworkInterface**
- **ec2>CreateSecurityGroup**
- **ec2:CreateTags**
- **ec2>CreateVolume**
- **ec2>DeleteSecurityGroup**
- **ec2>DeleteSnapshot**
- **ec2>DeleteTags**
- **ec2:DeregisterImage**
- **ec2:DescribeAccountAttributes**
- **ec2:DescribeAddresses**
- **ec2:DescribeAvailabilityZones**
- **ec2:DescribeDhcpOptions**
- **ec2:DescribeImages**
- **ec2:DescribeInstanceAttribute**
- **ec2:DescribeInstanceCreditSpecifications**
- **ec2:DescribeInstances**
- **ec2:DescribeInstanceTypes**
- **ec2:DescribeInternetGateways**

- **ec2:DescribeKeyPairs**
- **ec2:DescribeNatGateways**
- **ec2:DescribeNetworkAcls**
- **ec2:DescribeNetworkInterfaces**
- **ec2:DescribePrefixLists**
- **ec2:DescribePublicIpv4Pools** (only required if **publicIpv4Pool** is specified in **install-config.yaml**)
- **ec2:DescribeRegions**
- **ec2:DescribeRouteTables**
- **ec2:DescribeSecurityGroupRules**
- **ec2:DescribeSecurityGroups**
- **ec2:DescribeSubnets**
- **ec2:DescribeTags**
- **ec2:DescribeVolumes**
- **ec2:DescribeVpcAttribute**
- **ec2:DescribeVpcClassicLink**
- **ec2:DescribeVpcClassicLinkDnsSupport**
- **ec2:DescribeVpcEndpoints**
- **ec2:DescribeVpcs**
- **ec2:DisassociateAddress** (only required if **publicIpv4Pool** is specified in **install-config.yaml**)
- **ec2:GetEbsDefaultKmsKeyId**
- **ec2:ModifyInstanceAttribute**
- **ec2:ModifyNetworkInterfaceAttribute**
- **ec2:RevokeSecurityGroupEgress**
- **ec2:RevokeSecurityGroupIngress**
- **ec2:RunInstances**
- **ec2:TerminateInstances**

Example 6.58. Required permissions for creating network resources during installation

- **ec2:AllocateAddress**

- **ec2:AssociateAddress**
- **ec2:AssociateDhcpOptions**
- **ec2:AssociateRouteTable**
- **ec2:AttachInternetGateway**
- **ec2:CreateDhcpOptions**
- **ec2:CreateInternetGateway**
- **ec2:CreateNatGateway**
- **ec2:CreateRoute**
- **ec2:CreateRouteTable**
- **ec2:CreateSubnet**
- **ec2:CreateVpc**
- **ec2:CreateVpcEndpoint**
- **ec2:ModifySubnetAttribute**
- **ec2:ModifyVpcAttribute**

**NOTE**

If you use an existing Virtual Private Cloud (VPC), your account does not require these permissions for creating network resources.

Example 6.59. Required Elastic Load Balancing permissions (ELB) for installation

- **elasticloadbalancing:AddTags**
- **elasticloadbalancing:ApplySecurityGroupsToLoadBalancer**
- **elasticloadbalancing:AttachLoadBalancerToSubnets**
- **elasticloadbalancing:ConfigureHealthCheck**
- **elasticloadbalancing>CreateListener**
- **elasticloadbalancing>CreateLoadBalancer**
- **elasticloadbalancing>CreateLoadBalancerListeners**
- **elasticloadbalancing>CreateTargetGroup**
- **elasticloadbalancing>DeleteLoadBalancer**
- **elasticloadbalancing:DeregisterInstancesFromLoadBalancer**
- **elasticloadbalancing:DeregisterTargets**

- **elasticloadbalancing:DescribeInstanceHealth**
- **elasticloadbalancing:DescribeListeners**
- **elasticloadbalancing:DescribeLoadBalancerAttributes**
- **elasticloadbalancing:DescribeLoadBalancers**
- **elasticloadbalancing:DescribeTags**
- **elasticloadbalancing:DescribeTargetGroupAttributes**
- **elasticloadbalancing:DescribeTargetHealth**
- **elasticloadbalancing:ModifyLoadBalancerAttributes**
- **elasticloadbalancing:ModifyTargetGroup**
- **elasticloadbalancing:ModifyTargetGroupAttributes**
- **elasticloadbalancing:RegisterInstancesWithLoadBalancer**
- **elasticloadbalancing:RegisterTargets**
- **elasticloadbalancing:SetLoadBalancerPoliciesOfListener**
- **elasticloadbalancing:SetSecurityGroups**



IMPORTANT

OpenShift Container Platform uses both the ELB and ELBv2 API services to provision load balancers. The permission list shows permissions required by both services. A known issue exists in the AWS web console where both services use the same **elasticloadbalancing** action prefix but do not recognize the same actions. You can ignore the warnings about the service not recognizing certain **elasticloadbalancing** actions.

Example 6.60. Required IAM permissions for installation

- **iam:AddRoleToInstanceProfile**
- **iam:CreateInstanceProfile**
- **iam:CreateRole**
- **iam:DeleteInstanceProfile**
- **iam>DeleteRole**
- **iam>DeleteRolePolicy**
- **iam:GetInstanceProfile**
- **iam:GetRole**

- **iam:GetRolePolicy**
- **iam:GetUser**
- **iam:ListInstanceProfilesForRole**
- **iam:ListRoles**
- **iam:ListUsers**
- **iam:PassRole**
- **iam:PutRolePolicy**
- **iam:RemoveRoleFromInstanceProfile**
- **iam:SimulatePrincipalPolicy**
- **iam:TagInstanceProfile**
- **iam:TagRole**



NOTE

If you have not created a load balancer in your AWS account, the IAM user also requires the **iam:CreateServiceLinkedRole** permission.

Example 6.61. Required Route 53 permissions for installation

- **route53:ChangeResourceRecordSets**
- **route53:ChangeTagsForResource**
- **route53:CreateHostedZone**
- **route53>DeleteHostedZone**
- **route53:GetChange**
- **route53:GetHostedZone**
- **route53:ListHostedZones**
- **route53:ListHostedZonesByName**
- **route53:ListResourceRecordSets**
- **route53:ListTagsForResource**
- **route53:UpdateHostedZoneComment**

Example 6.62. Required Amazon Simple Storage Service (S3) permissions for installation

- **s3:CreateBucket**

- **s3:DeleteBucket**
- **s3:GetAccelerateConfiguration**
- **s3:GetBucketAcl**
- **s3:GetBucketCors**
- **s3:GetBucketLocation**
- **s3:GetBucketLogging**
- **s3:GetBucketObjectLockConfiguration**
- **s3:GetBucketPolicy**
- **s3:GetBucketRequestPayment**
- **s3:GetBucketTagging**
- **s3:GetBucketVersioning**
- **s3:GetBucketWebsite**
- **s3:GetEncryptionConfiguration**
- **s3:GetLifecycleConfiguration**
- **s3:GetReplicationConfiguration**
- **s3:ListBucket**
- **s3:PutBucketAcl**
- **s3:PutBucketPolicy**
- **s3:PutBucketTagging**
- **s3:PutEncryptionConfiguration**

Example 6.63. S3 permissions that cluster Operators require

- **s3:DeleteObject**
- **s3:GetObject**
- **s3:GetObjectAcl**
- **s3:GetObjectTagging**
- **s3:GetObjectVersion**
- **s3:PutObject**
- **s3:PutObjectAcl**
- **s3:PutObjectTagging**

Example 6.64. Required permissions to delete base cluster resources

- **autoscaling:DescribeAutoScalingGroups**
- **ec2:DeleteNetworkInterface**
- **ec2:DeletePlacementGroup**
- **ec2:DeleteVolume**
- **elasticloadbalancing:DeleteTargetGroup**
- **elasticloadbalancing:DescribeTargetGroups**
- **iam:DeleteAccessKey**
- **iam:DeleteUser**
- **iam:DeleteUserPolicy**
- **iam>ListAttachedRolePolicies**
- **iam>ListInstanceProfiles**
- **iam>ListRolePolicies**
- **iam>ListUserPolicies**
- **s3>DeleteObject**
- **s3>ListBucketVersions**
- **tag:GetResources**

Example 6.65. Required permissions to delete network resources

- **ec2:DeleteDhcpOptions**
- **ec2:DeleteInternetGateway**
- **ec2:DeleteNatGateway**
- **ec2:DeleteRoute**
- **ec2:DeleteRouteTable**
- **ec2:DeleteSubnet**
- **ec2:DeleteVpc**
- **ec2:DeleteVpcEndpoints**
- **ec2:DetachInternetGateway**

- **ec2:DisassociateRouteTable**
- **ec2:ReleaseAddress**
- **ec2:ReplaceRouteTableAssociation**

**NOTE**

If you use an existing VPC, your account does not require these permissions to delete network resources. Instead, your account only requires the **tag:UntagResources** permission to delete network resources.

Example 6.66. Optional permissions for installing a cluster with a custom Key Management Service (KMS) key

- **kms:CreateGrant**
- **kms:Decrypt**
- **kms:DescribeKey**
- **kms:Encrypt**
- **kms:GenerateDataKey**
- **kms:GenerateDataKeyWithoutPlainText**
- **kms:ListGrants**
- **kms:RevokeGrant**

Example 6.67. Required permissions to delete a cluster with shared instance roles

- **iam:UntagRole**

Example 6.68. Additional IAM and S3 permissions that are required to create manifests

- **iam:GetUserPolicy**
- **iam:ListAccessKeys**
- **iam:PutUserPolicy**
- **iam:TagUser**
- **s3:AbortMultipartUpload**
- **s3:GetBucketPublicAccessBlock**
- **s3:ListBucket**
- **s3:ListBucketMultipartUploads**

- **s3:PutBucketPublicAccessBlock**
- **s3:PutLifecycleConfiguration**



NOTE

If you are managing your cloud provider credentials with mint mode, the IAM user also requires the **iam:CreateAccessKey** and **iam:CreateUser** permissions.

Example 6.69. Optional permissions for instance and quota checks for installation

- **ec2:DescribeInstanceTypeOfferings**
- **servicequotas:ListAWSDefaultServiceQuotas**

Example 6.70. Optional permissions for the cluster owner account when installing a cluster on a shared VPC

- **sts:AssumeRole**

Example 6.71. Required permissions for enabling Bring your own public IPv4 addresses (BYOIP) feature for installation

- **ec2:DescribePublicIpv4Pools**
- **ec2:DisassociateAddress**

6.4.2.5. Obtaining an AWS Marketplace image

If you are deploying an OpenShift Container Platform cluster using an AWS Marketplace image, you must first subscribe through AWS. Subscribing to the offer provides you with the AMI ID that the installation program uses to deploy compute nodes.

Prerequisites

- You have an AWS account to purchase the offer. This account does not have to be the same account that is used to install the cluster.

Procedure

1. Complete the OpenShift Container Platform subscription from the [AWS Marketplace](#).

6.4.3. Installing a cluster on user-provisioned infrastructure in AWS by using CloudFormation templates

In OpenShift Container Platform version 4.16, you can install a cluster on Amazon Web Services (AWS) that uses infrastructure that you provide.

One way to create this infrastructure is to use the provided CloudFormation templates. You can modify the templates to customize your infrastructure or use the information that they contain to create AWS objects according to your company's policies.



IMPORTANT

The steps for performing a user-provisioned infrastructure installation are provided as an example only. Installing a cluster with infrastructure you provide requires knowledge of the cloud provider and the installation process of OpenShift Container Platform. Several CloudFormation templates are provided to assist in completing these steps or to help model your own. You are also free to create the required resources through other methods; the templates are just an example.

6.4.3.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [configured an AWS account](#) to host the cluster.



IMPORTANT

If you have an AWS profile stored on your computer, it must not use a temporary session token that you generated while using a multi-factor authentication device. The cluster continues to use your current AWS credentials to create AWS resources for the entire life of the cluster, so you must use key-based, long-term credentials. To generate appropriate keys, see [Managing Access Keys for IAM Users](#) in the AWS documentation. You can supply the keys when you run the installation program.

- You [prepared the user-provisioned infrastructure](#).
- You downloaded the AWS CLI and installed it on your computer. See [Install the AWS CLI Using the Bundled Installer \(Linux, macOS, or UNIX\)](#) in the AWS documentation.
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.



NOTE

Be sure to also review this site list if you are configuring a proxy.

- If the cloud identity and access management (IAM) APIs are not accessible in your environment, or if you do not want to store an administrator-level credential secret in the **kube-system** namespace, you can [manually create and maintain long-term credentials](#).

6.4.3.2. Creating the installation files for AWS

To install OpenShift Container Platform on Amazon Web Services (AWS) using user-provisioned infrastructure, you must generate the files that the installation program needs to deploy your cluster and modify them so that the cluster creates only the machines that it will use. You generate and

customize the **install-config.yaml** file, Kubernetes manifests, and Ignition config files. You also have the option to first set up a separate **var** partition during the preparation phases of installation.

6.4.3.2.1. Optional: Creating a separate **/var** partition


It is recommended that disk partitioning for OpenShift Container Platform be left to the installer. However, there are cases where you might want to create separate partitions in a part of the filesystem that you expect to grow.

OpenShift Container Platform supports the addition of a single partition to attach storage to either the **/var** partition or a subdirectory of **/var**. For example:

- **/var/lib/containers**: Holds container-related content that can grow as more images and containers are added to a system.
- **/var/lib/etcd**: Holds data that you might want to keep separate for purposes such as performance optimization of etcd storage.
- **/var**: Holds data that you might want to keep separate for purposes such as auditing.

Storing the contents of a **/var** directory separately makes it easier to grow storage for those areas as needed and reinstall OpenShift Container Platform at a later date and keep that data intact. With this method, you will not have to pull all your containers again, nor will you have to copy massive log files when you update systems.

Because **/var** must be in place before a fresh installation of Red Hat Enterprise Linux CoreOS (RHCOS), the following procedure sets up the separate **/var** partition by creating a machine config manifest that is inserted during the **openshift-install** preparation phases of an OpenShift Container Platform installation.



IMPORTANT

If you follow the steps to create a separate **/var** partition in this procedure, it is not necessary to create the Kubernetes manifest and Ignition config files again as described later in this section.

Procedure

1. Create a directory to hold the OpenShift Container Platform installation files:

```
$ mkdir $HOME/clusterconfig
```

2. Run **openshift-install** to create a set of files in the **manifest** and **openshift** subdirectories. Answer the system questions as you are prompted:

```
$ openshift-install create manifests --dir $HOME/clusterconfig
```

Example output

```
? SSH Public Key ...
INFO Credentials loaded from the "myprofile" profile in file "/home/myuser/.aws/credentials"
INFO Consuming Install Config from target directory
INFO Manifests created in: $HOME/clusterconfig/manifests and
$HOME/clusterconfig/openshift
```

- Optional: Confirm that the installation program created manifests in the **clusterconfig/openshift** directory:

```
$ ls $HOME/clusterconfig/openshift/
```


Example output

```
99_kubeadmin-password-secret.yaml
99_openshift-cluster-api_master-machines-0.yaml
99_openshift-cluster-api_master-machines-1.yaml
99_openshift-cluster-api_master-machines-2.yaml
...
```

- Create a Butane config that configures the additional partition. For example, name the file **\$HOME/clusterconfig/98-var-partition.bu**, change the disk device name to the name of the storage device on the **worker** systems, and set the storage size as appropriate. This example places the **/var** directory on a separate partition:

```
variant: openshift
version: 4.16.0
metadata:
  labels:
    machineconfiguration.openshift.io/role: worker
  name: 98-var-partition
storage:
  disks:
    - device: /dev/disk/by-id/<device_name> 1
      partitions:
        - label: var
          start_mib: <partition_start_offset> 2
          size_mib: <partition_size> 3
          number: 5
      filesystems:
        - device: /dev/disk/by-partlabel/var
          path: /var
          format: xfs
          mount_options: [defaults, prjquota] 4
          with_mount_unit: true
```

- The storage device name of the disk that you want to partition.
- When adding a data partition to the boot disk, a minimum value of 25000 MiB (Mebibytes) is recommended. The root file system is automatically resized to fill all available space up to the specified offset. If no value is specified, or if the specified value is smaller than the recommended minimum, the resulting root file system will be too small, and future reinstalls of RHCOS might overwrite the beginning of the data partition.
- The size of the data partition in mebibytes.
- The **prjquota** mount option must be enabled for filesystems used for container storage.

**NOTE**

When creating a separate **/var** partition, you cannot use different instance types for worker nodes, if the different instance types do not have the same device name.

5. Create a manifest from the Butane config and save it to the **clusterconfig/openshift** directory. For example, run the following command:

```
$ butane $HOME/clusterconfig/98-var-partition.bu -o $HOME/clusterconfig/openshift/98-var-partition.yaml
```

6. Run **openshift-install** again to create Ignition configs from a set of files in the **manifest** and **openshift** subdirectories:

```
$ openshift-install create ignition-configs --dir $HOME/clusterconfig
$ ls $HOME/clusterconfig/
auth bootstrap.ign master.ign metadata.json worker.ign
```

Now you can use the Ignition config files as input to the installation procedures to install Red Hat Enterprise Linux CoreOS (RHCOS) systems.

6.4.3.2.2. Creating the installation configuration file

Generate and customize the installation configuration file that the installation program needs to deploy your cluster.

Prerequisites

- You obtained the OpenShift Container Platform installation program for user-provisioned infrastructure and the pull secret for your cluster.
- You checked that you are deploying your cluster to an AWS Region with an accompanying Red Hat Enterprise Linux CoreOS (RHCOS) AMI published by Red Hat. If you are deploying to an AWS Region that requires a custom AMI, such as an AWS GovCloud Region, you must create the **install-config.yaml** file manually.

Procedure

1. Create the **install-config.yaml** file.
 - a. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install create install-config --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the directory name to store the files that the installation program creates.

**IMPORTANT**

Specify an empty directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

- b. At the prompts, provide the configuration details for your cloud:
 - i. Optional: Select an SSH key to use to access your cluster machines.

**NOTE**

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- ii. Select **aws** as the platform to target.
- iii. If you do not have an AWS profile stored on your computer, enter the AWS access key ID and secret access key for the user that you configured to run the installation program.

**NOTE**

The AWS access key ID and secret access key are stored in `~/.aws/credentials` in the home directory of the current user on the installation host. You are prompted for the credentials by the installation program if the credentials for the exported profile are not present in the file. Any credentials that you provide to the installation program are stored in the file.

- iv. Select the AWS Region to deploy the cluster to.
 - v. Select the base domain for the Route 53 service that you configured for your cluster.
 - vi. Enter a descriptive name for your cluster.
 - vii. Paste the [pull secret from Red Hat OpenShift Cluster Manager](#) .
2. If you are installing a three-node cluster, modify the **install-config.yaml** file by setting the **compute.replicas** parameter to **0**. This ensures that the cluster's control planes are schedulable. For more information, see "Installing a three-node cluster on AWS".
 3. Optional: Back up the **install-config.yaml** file.

**IMPORTANT**

The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

Additional resources


- See [Configuration and credential file settings](#) in the AWS documentation for more information about AWS profile and credential configuration.

6.4.3.2.3. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: ec2.<aws_region>.amazonaws.com,elasticloadbalancing.
  <aws_region>.amazonaws.com,s3.<aws_region>.amazonaws.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.

2 A proxy URL to use for creating HTTPS connections outside the cluster.

3

A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with `.` to match subdomains only. For

- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.



NOTE

The installation program does not support the proxy **readinessEndpoints** field.



NOTE

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.



NOTE

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

6.4.3.2.4. Creating the Kubernetes manifest and Ignition config files

Because you must modify some cluster definition files and manually start the cluster machines, you must generate the Kubernetes manifest and Ignition config files that the cluster needs to configure the machines.

The installation configuration file transforms into the Kubernetes manifests. The manifests wrap into the Ignition configuration files, which are later used to configure the cluster machines.



IMPORTANT

- The Ignition config files that the OpenShift Container Platform installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

Prerequisites

- You obtained the OpenShift Container Platform installation program.
- You created the **install-config.yaml** installation configuration file.

Procedure

1. Change to the directory that contains the OpenShift Container Platform installation program and generate the Kubernetes manifests for the cluster:

```
$ ./openshift-install create manifests --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the installation directory that contains the **install-config.yaml** file you created.

2. Remove the Kubernetes manifest files that define the control plane machines:


```
$ rm -f <installation_directory>/openshift/99_openshift-cluster-api_master-machines-*.yaml
```

By removing these files, you prevent the cluster from automatically generating control plane machines.

3. Remove the Kubernetes manifest files that define the control plane machine set:

```
$ rm -f <installation_directory>/openshift/99_openshift-machine-api_master-control-plane-machine-set.yaml
```

```
$ rm -f <installation_directory>/openshift/99_openshift-cluster-api_worker-machineset-*.yaml
```



IMPORTANT


If you disabled the **MachineAPI** capability when installing a cluster on user-provisioned infrastructure, you must remove the Kubernetes manifest files that define the worker machines. Otherwise, your cluster fails to install.

Because you create and manage the worker machines yourself, you do not need to initialize these machines.



WARNING

If you are installing a three-node cluster, skip the following step to allow the control plane nodes to be schedulable.



IMPORTANT

When you configure control plane nodes from the default unschedulable to schedulable, additional subscriptions are required. This is because control plane nodes then become compute nodes.

4. Check that the **mastersSchedulable** parameter in the `<installation_directory>/manifests/cluster-scheduler-02-config.yml` Kubernetes manifest file is set to **false**. This setting prevents pods from being scheduled on the control plane machines:
 - a. Open the `<installation_directory>/manifests/cluster-scheduler-02-config.yml` file.
 - b. Locate the **mastersSchedulable** parameter and ensure that it is set to **false**.
 - c. Save and exit the file.
5. Optional: If you do not want [the Ingress Operator](#) to create DNS records on your behalf, remove the **privateZone** and **publicZone** sections from the `<installation_directory>/manifests/cluster-dns-02-config.yml` DNS configuration file:

```
apiVersion: config.openshift.io/v1
kind: DNS
metadata:
  creationTimestamp: null
  name: cluster
spec:
  baseDomain: example.openshift.com
  privateZone: 1
    id: mycluster-100419-private-zone
  publicZone: 2
    id: example.openshift.com
status: {}
```

- 1 2 Remove this section completely.


If you do so, you must add ingress DNS records manually in a later step.

6. To create the Ignition configuration files, run the following command from the directory that contains the installation program:

```
$ ./openshift-install create ignition-configs --dir <installation_directory> 1
```

- 1 For `<installation_directory>`, specify the same installation directory.

Ignition config files are created for the bootstrap, control plane, and compute nodes in the installation directory. The `kubeadmin-password` and `kubeconfig` files are created in the `./<installation_directory>/auth` directory:



6.4.3.3. Extracting the infrastructure name

The Ignition config files contain a unique cluster identifier that you can use to uniquely identify your cluster in Amazon Web Services (AWS). The infrastructure name is also used to locate the appropriate AWS resources during an OpenShift Container Platform installation. The provided CloudFormation templates contain references to this infrastructure name, so you must extract it.

Prerequisites

- You obtained the OpenShift Container Platform installation program and the pull secret for your cluster.
- You generated the Ignition config files for your cluster.
- You installed the `jq` package.

Procedure

- To extract and view the infrastructure name from the Ignition config file metadata, run the following command:

```
$ jq -r .infraID <installation_directory>/metadata.json 1
```

- 1 For `<installation_directory>`, specify the path to the directory that you stored the installation files in.

Example output

```
openshift-vw9j6 1
```

- 1 The output of this command is your cluster name and a random string.

6.4.3.4. Creating a VPC in AWS

You must create a Virtual Private Cloud (VPC) in Amazon Web Services (AWS) for your OpenShift Container Platform cluster to use. You can customize the VPC to meet your requirements, including VPN and route tables.

You can use the provided CloudFormation template and a custom parameter file to create a stack of AWS resources that represent the VPC.



NOTE

If you do not use the provided CloudFormation template to create your AWS infrastructure, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- You configured an AWS account.
- You added your AWS keys and region to your local AWS profile by running **aws configure**.
- You generated the Ignition config files for your cluster.

Procedure

1. Create a JSON file that contains the parameter values that the template requires:

```
[
  {
    "ParameterKey": "VpcCidr", 1
    "ParameterValue": "10.0.0.0/16" 2
  },
  {
    "ParameterKey": "AvailabilityZoneCount", 3
    "ParameterValue": "1" 4
  },
  {
    "ParameterKey": "SubnetBits", 5
    "ParameterValue": "12" 6
  }
]
```

- 1** The CIDR block for the VPC.
- 2** Specify a CIDR block in the format **x.x.x.x/16-24**.
- 3** The number of availability zones to deploy the VPC in.
- 4** Specify an integer between **1** and **3**.
- 5** The size of each subnet in each availability zone.
- 6** Specify an integer between **5** and **13**, where **5** is **/27** and **13** is **/19**.

2. Copy the template from the **CloudFormation template for the VPC** section of this topic and save it as a YAML file on your computer. This template describes the VPC that your cluster requires.
3. Launch the CloudFormation template to create a stack of AWS resources that represent the VPC:



IMPORTANT

You must enter the command on a single line.

```
$ aws cloudformation create-stack --stack-name <name> 1
  --template-body file://<template>.yaml 2
  --parameters file://<parameters>.json 3
```

- 1** **<name>** is the name for the CloudFormation stack, such as **cluster-vpc**. You need the name of this stack if you remove the cluster.
- 2** **<template>** is the relative path to and name of the CloudFormation template YAML file that you saved.
- 3** **<parameters>** is the relative path to and name of the CloudFormation parameters JSON file.

Example output

```
arn:aws:cloudformation:us-east-1:269333783861:stack/cluster-vpc/dbedae40-2fd3-11eb-820e-12a48460849f
```

4. Confirm that the template components exist:

```
$ aws cloudformation describe-stacks --stack-name <name>
```

After the **StackStatus** displays **CREATE_COMPLETE**, the output displays values for the following parameters. You must provide these parameter values to the other CloudFormation templates that you run to create your cluster:

VpcId	The ID of your VPC.
PublicSubnetIds	The IDs of the new public subnets.
PrivateSubnetIds	The IDs of the new private subnets.

6.4.3.4.1. CloudFormation template for the VPC

You can use the following CloudFormation template to deploy the VPC that you need for your OpenShift Container Platform cluster.

Example 6.72. CloudFormation template for the VPC

AWSTemplateFormatVersion: 2010-09-09

Description: Template for Best Practice VPC with 1-3 AZs

Parameters:

VpcCidr:

AllowedPattern: ^(((0-9){1-9}|0-9|1[0-9]{2}|2[0-4][0-9]|25[0-5])\.)\{3\}([0-9]{1-9}|0-9|1[0-9]{2}|2[0-4][0-9]|25[0-5])\(\{1[6-9]|2[0-4]\}\)\$

ConstraintDescription: CIDR block parameter must be in the form x.x.x.x/16-24.

Default: 10.0.0.0/16

Description: CIDR block for VPC.

Type: String

AvailabilityZoneCount:

ConstraintDescription: "The number of availability zones. (Min: 1, Max: 3)"

MinValue: 1

MaxValue: 3

Default: 1

Description: "How many AZs to create VPC subnets for. (Min: 1, Max: 3)"

Type: Number

SubnetBits:

ConstraintDescription: CIDR block parameter must be in the form x.x.x.x/19-27.

MinValue: 5

MaxValue: 13

Default: 12

Description: "Size of each subnet to create within the availability zones. (Min: 5 = /27, Max: 13 = /19)"

Type: Number

Metadata:

AWS::CloudFormation::Interface:

ParameterGroups:

- Label:

default: "Network Configuration"

Parameters:

- VpcCidr

- SubnetBits

- Label:

default: "Availability Zones"

Parameters:

- AvailabilityZoneCount

ParameterLabels:

AvailabilityZoneCount:

default: "Availability Zone Count"

VpcCidr:

default: "VPC CIDR"

SubnetBits:

default: "Bits Per Subnet"

Conditions:

DoAz3: !Equals [3, !Ref AvailabilityZoneCount]

DoAz2: !Or [!Equals [2, !Ref AvailabilityZoneCount], Condition: DoAz3]

Resources:

VPC:

Type: "AWS::EC2::VPC"

Properties:

EnableDnsSupport: "true"


```

    EnableDnsHostnames: "true"
    CidrBlock: !Ref VpcCidr
PublicSubnet:
  Type: "AWS::EC2::Subnet"
  Properties:
    VpId: !Ref VPC
    CidrBlock: !Select [0, !Cidr [!Ref VpcCidr, 6, !Ref SubnetBits]]
    AvailabilityZone: !Select
      - 0
      - Fn::GetAZs: !Ref "AWS::Region"
PublicSubnet2:
  Type: "AWS::EC2::Subnet"
  Condition: DoAz2
  Properties:
    VpId: !Ref VPC
    CidrBlock: !Select [1, !Cidr [!Ref VpcCidr, 6, !Ref SubnetBits]]
    AvailabilityZone: !Select
      - 1
      - Fn::GetAZs: !Ref "AWS::Region"
PublicSubnet3:
  Type: "AWS::EC2::Subnet"
  Condition: DoAz3
  Properties:
    VpId: !Ref VPC
    CidrBlock: !Select [2, !Cidr [!Ref VpcCidr, 6, !Ref SubnetBits]]
    AvailabilityZone: !Select
      - 2
      - Fn::GetAZs: !Ref "AWS::Region"
InternetGateway:
  Type: "AWS::EC2::InternetGateway"
GatewayToInternet:
  Type: "AWS::EC2::VPCGatewayAttachment"
  Properties:
    VpId: !Ref VPC
    InternetGatewayId: !Ref InternetGateway
PublicRouteTable:
  Type: "AWS::EC2::RouteTable"
  Properties:
    VpId: !Ref VPC
PublicRoute:
  Type: "AWS::EC2::Route"
  DependsOn: GatewayToInternet
  Properties:
    RouteTableId: !Ref PublicRouteTable
    DestinationCidrBlock: 0.0.0.0/0
    GatewayId: !Ref InternetGateway
PublicSubnetRouteTableAssociation:
  Type: "AWS::EC2::SubnetRouteTableAssociation"
  Properties:
    SubnetId: !Ref PublicSubnet
    RouteTableId: !Ref PublicRouteTable
PublicSubnetRouteTableAssociation2:
  Type: "AWS::EC2::SubnetRouteTableAssociation"
  Condition: DoAz2
  Properties:
    SubnetId: !Ref PublicSubnet2

```

```

RouteTableId: !Ref PublicRouteTable
PublicSubnetRouteTableAssociation3:
  Condition: DoAz3
  Type: "AWS::EC2::SubnetRouteTableAssociation"
  Properties:
    SubnetId: !Ref PublicSubnet3
    RouteTableId: !Ref PublicRouteTable
PrivateSubnet:
  Type: "AWS::EC2::Subnet"
  Properties:
    VpcId: !Ref VPC
    CidrBlock: !Select [3, !Cidr [!Ref VpcCidr, 6, !Ref SubnetBits]]
    AvailabilityZone: !Select
      - 0
      - Fn::GetAZs: !Ref "AWS::Region"
PrivateRouteTable:
  Type: "AWS::EC2::RouteTable"
  Properties:
    VpcId: !Ref VPC
PrivateSubnetRouteTableAssociation:
  Type: "AWS::EC2::SubnetRouteTableAssociation"
  Properties:
    SubnetId: !Ref PrivateSubnet
    RouteTableId: !Ref PrivateRouteTable
NAT:
  DependsOn:
    - GatewayToInternet
  Type: "AWS::EC2::NatGateway"
  Properties:
    AllocationId:
      "Fn::GetAtt":
        - EIP
        - AllocationId
    SubnetId: !Ref PublicSubnet
EIP:
  Type: "AWS::EC2::EIP"
  Properties:
    Domain: vpc
Route:
  Type: "AWS::EC2::Route"
  Properties:
    RouteTableId:
      Ref: PrivateRouteTable
    DestinationCidrBlock: 0.0.0.0/0
    NatGatewayId:
      Ref: NAT
PrivateSubnet2:
  Type: "AWS::EC2::Subnet"
  Condition: DoAz2
  Properties:
    VpcId: !Ref VPC
    CidrBlock: !Select [4, !Cidr [!Ref VpcCidr, 6, !Ref SubnetBits]]
    AvailabilityZone: !Select
      - 1
      - Fn::GetAZs: !Ref "AWS::Region"
PrivateRouteTable2:

```

```

Type: "AWS::EC2::RouteTable"
Condition: DoAz2
Properties:
  VpcId: !Ref VPC
PrivateSubnetRouteTableAssociation2:
  Type: "AWS::EC2::SubnetRouteTableAssociation"
  Condition: DoAz2
  Properties:
    SubnetId: !Ref PrivateSubnet2
    RouteTableId: !Ref PrivateRouteTable2
NAT2:
  DependsOn:
  - GatewayToInternet
  Type: "AWS::EC2::NatGateway"
  Condition: DoAz2
  Properties:
    AllocationId:
      "Fn::GetAtt":
      - EIP2
      - AllocationId
    SubnetId: !Ref PublicSubnet2
EIP2:
  Type: "AWS::EC2::EIP"
  Condition: DoAz2
  Properties:
    Domain: vpc
Route2:
  Type: "AWS::EC2::Route"
  Condition: DoAz2
  Properties:
    RouteTableId:
      Ref: PrivateRouteTable2
    DestinationCidrBlock: 0.0.0.0/0
    NatGatewayId:
      Ref: NAT2
PrivateSubnet3:
  Type: "AWS::EC2::Subnet"
  Condition: DoAz3
  Properties:
    VpcId: !Ref VPC
    CidrBlock: !Select [5, !Cidr [!Ref VpcCidr, 6, !Ref SubnetBits]]
    AvailabilityZone: !Select
      - 2
      - Fn::GetAZs: !Ref "AWS::Region"
PrivateRouteTable3:
  Type: "AWS::EC2::RouteTable"
  Condition: DoAz3
  Properties:
    VpcId: !Ref VPC
PrivateSubnetRouteTableAssociation3:
  Type: "AWS::EC2::SubnetRouteTableAssociation"
  Condition: DoAz3
  Properties:
    SubnetId: !Ref PrivateSubnet3
    RouteTableId: !Ref PrivateRouteTable3
NAT3:

```

```

DependsOn:
- GatewayToInternet
Type: "AWS::EC2::NatGateway"
Condition: DoAz3
Properties:
  AllocationId:
    "Fn::GetAtt":
      - EIP3
      - AllocationId
  SubnetId: !Ref PublicSubnet3
EIP3:
  Type: "AWS::EC2::EIP"
  Condition: DoAz3
  Properties:
    Domain: vpc
Route3:
  Type: "AWS::EC2::Route"
  Condition: DoAz3
  Properties:
    RouteTableId:
      Ref: PrivateRouteTable3
    DestinationCidrBlock: 0.0.0.0/0
    NatGatewayId:
      Ref: NAT3
S3Endpoint:
  Type: AWS::EC2::VPCEndpoint
  Properties:
    PolicyDocument:
      Version: 2012-10-17
      Statement:
        - Effect: Allow
          Principal: '*'
          Action:
            - '*'
          Resource:
            - '*'
    RouteTableIds:
      - !Ref PublicRouteTable
      - !Ref PrivateRouteTable
      - !If [DoAz2, !Ref PrivateRouteTable2, !Ref "AWS::NoValue"]
      - !If [DoAz3, !Ref PrivateRouteTable3, !Ref "AWS::NoValue"]
    ServiceName: !Join
      - "
      - - com.amazonaws.
      - - !Ref 'AWS::Region'
      - - .s3
    Vpclid: !Ref VPC

Outputs:
Vpclid:
  Description: ID of the new VPC.
  Value: !Ref VPC
PublicSubnetIds:
  Description: Subnet IDs of the public subnets.
  Value:
    !Join [

```

```

    ""
    [!Ref PublicSubnet, !If [DoAz2, !Ref PublicSubnet2, !Ref "AWS::NoValue"], !If [DoAz3, !Ref
PublicSubnet3, !Ref "AWS::NoValue"]]
  ]
  PrivateSubnetIds:
    Description: Subnet IDs of the private subnets.
    Value:
      !Join [
        "",
        [!Ref PrivateSubnet, !If [DoAz2, !Ref PrivateSubnet2, !Ref "AWS::NoValue"], !If [DoAz3, !Ref
PrivateSubnet3, !Ref "AWS::NoValue"]]
      ]
  PublicRouteTableId:
    Description: Public Route table ID
    Value: !Ref PublicRouteTable
  PrivateRouteTableIds:
    Description: Private Route table IDs
    Value:
      !Join [
        "",
        [
          !Join ["=", [
            !Select [0, "Fn::GetAZs": !Ref "AWS::Region"],
            !Ref PrivateRouteTable
          ]],
          !If [DoAz2,
            !Join ["=", [!Select [1, "Fn::GetAZs": !Ref "AWS::Region"], !Ref PrivateRouteTable2]],
            !Ref "AWS::NoValue"
          ],
          !If [DoAz3,
            !Join ["=", [!Select [2, "Fn::GetAZs": !Ref "AWS::Region"], !Ref PrivateRouteTable3]],
            !Ref "AWS::NoValue"
          ]
        ]
      ]
  ]
]

```

Additional resources

- You can view details about the CloudFormation stacks that you create by navigating to the [AWS CloudFormation console](#).

6.4.3.5. Creating networking and load balancing components in AWS

You must configure networking and classic or network load balancing in Amazon Web Services (AWS) that your OpenShift Container Platform cluster can use.

You can use the provided CloudFormation template and a custom parameter file to create a stack of AWS resources. The stack represents the networking and load balancing components that your OpenShift Container Platform cluster requires. The template also creates a hosted zone and subnet tags.

You can run the template multiple times within a single Virtual Private Cloud (VPC).



NOTE

If you do not use the provided CloudFormation template to create your AWS infrastructure, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- You configured an AWS account.
- You added your AWS keys and region to your local AWS profile by running **aws configure**.
- You generated the Ignition config files for your cluster.
- You created and configured a VPC and associated subnets in AWS.

Procedure

1. Obtain the hosted zone ID for the Route 53 base domain that you specified in the **install-config.yaml** file for your cluster. You can obtain details about your hosted zone by running the following command:

```
$ aws route53 list-hosted-zones-by-name --dns-name <route53_domain> 1
```

- 1 For the **<route53_domain>**, specify the Route 53 base domain that you used when you generated the **install-config.yaml** file for the cluster.

Example output

```
mycluster.example.com. False 100
HOSTEDZONES 65F8F38E-2268-B835-E15C-AB55336FCBFA
/hostedzone/Z21IXYZABCZ2A4 mycluster.example.com. 10
```

In the example output, the hosted zone ID is **Z21IXYZABCZ2A4**.

2. Create a JSON file that contains the parameter values that the template requires:


```
[
  {
    "ParameterKey": "ClusterName", 1
    "ParameterValue": "mycluster" 2
  },
  {
    "ParameterKey": "InfrastructureName", 3
    "ParameterValue": "mycluster-<random_string>" 4
  },
  {
    "ParameterKey": "HostedZoneId", 5
    "ParameterValue": "<random_string>" 6
  },
  {
    "ParameterKey": "HostedZoneName", 7
  }
]
```

```

    "ParameterValue": "example.com" 8
  },
  {
    "ParameterKey": "PublicSubnets", 9
    "ParameterValue": "subnet-<random_string>" 10
  },
  {
    "ParameterKey": "PrivateSubnets", 11
    "ParameterValue": "subnet-<random_string>" 12
  },
  {
    "ParameterKey": "VpcId", 13
    "ParameterValue": "vpc-<random_string>" 14
  }
]


```

- 1 A short, representative cluster name to use for hostnames, etc.
 - 2 Specify the cluster name that you used when you generated the **install-config.yaml** file for the cluster.
 - 3 The name for your cluster infrastructure that is encoded in your Ignition config files for the cluster.
 - 4 Specify the infrastructure name that you extracted from the Ignition config file metadata, which has the format **<cluster-name>-<random-string>**.
 - 5 The Route 53 public zone ID to register the targets with.
 - 6 Specify the Route 53 public zone ID, which as a format similar to **Z21IXYZABCZ2A4**. You can obtain this value from the AWS console.
 - 7 The Route 53 zone to register the targets with.
 - 8 Specify the Route 53 base domain that you used when you generated the **install-config.yaml** file for the cluster. Do not include the trailing period (.) that is displayed in the AWS console.
 - 9 The public subnets that you created for your VPC.
 - 10 Specify the **PublicSubnetIds** value from the output of the CloudFormation template for the VPC.
 - 11 The private subnets that you created for your VPC.
 - 12 Specify the **PrivateSubnetIds** value from the output of the CloudFormation template for the VPC.
 - 13 The VPC that you created for the cluster.
 - 14 Specify the **VpcId** value from the output of the CloudFormation template for the VPC.
3. Copy the template from the **CloudFormation template for the network and load balancers** section of this topic and save it as a YAML file on your computer. This template describes the networking and load balancing objects that your cluster requires.

**IMPORTANT**

If you are deploying your cluster to an AWS government or secret region, you must update the **InternalApiServerRecord** in the CloudFormation template to use **CNAME** records. Records of type **ALIAS** are not supported for AWS government regions.

4. Launch the CloudFormation template to create a stack of AWS resources that provide the networking and load balancing components:

**IMPORTANT**

You must enter the command on a single line.

```
$ aws cloudformation create-stack --stack-name <name> 1
  --template-body file://<template>.yaml 2
  --parameters file://<parameters>.json 3
  --capabilities CAPABILITY_NAMED_IAM 4
```

- 1** **<name>** is the name for the CloudFormation stack, such as **cluster-dns**. You need the name of this stack if you remove the cluster.
- 2** **<template>** is the relative path to and name of the CloudFormation template YAML file that you saved.
- 3** **<parameters>** is the relative path to and name of the CloudFormation parameters JSON file.
- 4** You must explicitly declare the **CAPABILITY_NAMED_IAM** capability because the provided template creates some **AWS::IAM::Role** resources.

Example output

```
arn:aws:cloudformation:us-east-1:269333783861:stack/cluster-dns/cd3e5de0-2fd4-11eb-5cf0-12be5c33a183
```

5. Confirm that the template components exist:

```
$ aws cloudformation describe-stacks --stack-name <name>
```

After the **StackStatus** displays **CREATE_COMPLETE**, the output displays values for the following parameters. You must provide these parameter values to the other CloudFormation templates that you run to create your cluster:

PrivateHostedZoneId	Hosted zone ID for the private DNS.
----------------------------	-------------------------------------

ExternalApiLoadBalancerName	Full name of the external API load balancer.
InternalApiLoadBalancerName	Full name of the internal API load balancer.
ApiServerDnsName	Full hostname of the API server.
RegisterNlbIpTargetsLambda	Lambda ARN useful to help register/deregister IP targets for these load balancers.
ExternalApiTargetGroupArn	ARN of external API target group.
InternalApiTargetGroupArn	ARN of internal API target group.
InternalServiceTargetGroupArn	ARN of internal service target group.

6.4.3.5.1. CloudFormation template for the network and load balancers

You can use the following CloudFormation template to deploy the networking objects and load balancers that you need for your OpenShift Container Platform cluster.

Example 6.73. CloudFormation template for the network and load balancers

AWSTemplateFormatVersion: 2010-09-09

Description: Template for OpenShift Cluster Network Elements (Route53 & LBs)

Parameters:

ClusterName:

AllowedPattern: ^([a-zA-Z][a-zA-Z0-9\-\-]{0,26})\$

MaxLength: 27

MinLength: 1

ConstraintDescription: Cluster name must be alphanumeric, start with a letter, and have a maximum of 27 characters.

Description: A short, representative cluster name to use for host names and other identifying names.

Type: String

InfrastructureName:

AllowedPattern: ^([a-zA-Z][a-zA-Z0-9\-\-]{0,26})\$

MaxLength: 27

MinLength: 1

ConstraintDescription: Infrastructure name must be alphanumeric, start with a letter, and have a maximum of 27 characters.

Description: A short, unique cluster ID used to tag cloud resources and identify items owned or used by the cluster.

Type: String

HostedZoneld:

Description: The Route53 public zone ID to register the targets with, such as Z21IXYZABCZ2A4.

Type: String

HostedZoneName:

Description: The Route53 zone to register the targets with, such as example.com. Omit the trailing period.

Type: String

Default: "example.com"

PublicSubnets:

Description: The internet-facing subnets.

Type: List<AWS::EC2::Subnet::Id>

PrivateSubnets:

Description: The internal subnets.

Type: List<AWS::EC2::Subnet::Id>

VpcId:

Description: The VPC-scoped resources will belong to this VPC.

Type: AWS::EC2::VPC::Id

Metadata:

AWS::CloudFormation::Interface:

ParameterGroups:

- Label:

default: "Cluster Information"

Parameters:

- ClusterName

- InfrastructureName

- Label:

default: "Network Configuration"

Parameters:

- VpcId

- PublicSubnets

- PrivateSubnets

- Label:

default: "DNS"

Parameters:

- HostedZoneName

- HostedZoneld

ParameterLabels:

ClusterName:

default: "Cluster Name"

InfrastructureName:

default: "Infrastructure Name"

VpcId:

default: "VPC ID"

PublicSubnets:

default: "Public Subnets"

PrivateSubnets:

default: "Private Subnets"

HostedZoneName:
 default: "Public Hosted Zone Name"
 HostedZoneId:
 default: "Public Hosted Zone ID"

Resources:

ExtApiElb:

Type: AWS::ElasticLoadBalancingV2::LoadBalancer
 Properties:
 Name: !Join ["-", [!Ref InfrastructureName, "ext"]]
 IpAddressType: ipv4
 Subnets: !Ref PublicSubnets
 Type: network

IntApiElb:

Type: AWS::ElasticLoadBalancingV2::LoadBalancer
 Properties:
 Name: !Join ["-", [!Ref InfrastructureName, "int"]]
 Scheme: internal
 IpAddressType: ipv4
 Subnets: !Ref PrivateSubnets
 Type: network

IntDns:

Type: "AWS::Route53::HostedZone"
 Properties:
 HostedZoneConfig:
 Comment: "Managed by CloudFormation"
 Name: !Join [".", [!Ref ClusterName, !Ref HostedZoneName]]
 HostedZoneTags:
 - Key: Name
 Value: !Join ["-", [!Ref InfrastructureName, "int"]]
 - Key: !Join ["", ["kubernetes.io/cluster/", !Ref InfrastructureName]]
 Value: "owned"
 VPCs:
 - VPCId: !Ref VpcId
 VPCRegion: !Ref "AWS::Region"

ExternalApiServerRecord:

Type: AWS::Route53::RecordSetGroup
 Properties:
 Comment: Alias record for the API server
 HostedZoneId: !Ref HostedZoneId
 RecordSets:
 - Name:
 !Join [
 ".",
 ["api", !Ref ClusterName, !Join ["", [!Ref HostedZoneName, "."]]],
]
 Type: A
 AliasTarget:
 HostedZoneId: !GetAtt ExtApiElb.CanonicalHostedZoneID
 DNSName: !GetAtt ExtApiElb.DNSName

InternalApiServerRecord:

Type: AWS::Route53::RecordSetGroup

Properties:

Comment: Alias record for the API server

HostedZoneId: !Ref IntDns

RecordSets:

- Name:

```
!Join [
  ".",
  ["api", !Ref ClusterName, !Join ["", [!Ref HostedZoneName, "."]]],
]
```

Type: A

AliasTarget:

HostedZoneId: !GetAtt IntApiElb.CanonicalHostedZoneID

DNSName: !GetAtt IntApiElb.DNSName

- Name:

```
!Join [
  ".",
  ["api-int", !Ref ClusterName, !Join ["", [!Ref HostedZoneName, "."]]],
]
```

Type: A

AliasTarget:

HostedZoneId: !GetAtt IntApiElb.CanonicalHostedZoneID

DNSName: !GetAtt IntApiElb.DNSName

ExternalApiListener:

Type: AWS::ElasticLoadBalancingV2::Listener

Properties:

DefaultActions:

- Type: forward

TargetGroupArn:

Ref: ExternalApiTargetGroup

LoadBalancerArn:

Ref: ExtApiElb

Port: 6443

Protocol: TCP

ExternalApiTargetGroup:

Type: AWS::ElasticLoadBalancingV2::TargetGroup

Properties:

HealthCheckIntervalSeconds: 10

HealthCheckPath: "/readyz"

HealthCheckPort: 6443

HealthCheckProtocol: HTTPS

HealthyThresholdCount: 2

UnhealthyThresholdCount: 2

Port: 6443

Protocol: TCP

TargetType: ip

VpcId:

Ref: VpcId

TargetGroupAttributes:

- Key: deregistration_delay.timeout_seconds

Value: 60

InternalApiListener:

Type: AWS::ElasticLoadBalancingV2::Listener

Properties:

DefaultActions:
- Type: forward
 TargetGroupArn:
 Ref: InternalApiTargetGroup
LoadBalancerArn:
 Ref: IntApiElb
Port: 6443
Protocol: TCP

InternalApiTargetGroup:
Type: AWS::ElasticLoadBalancingV2::TargetGroup
Properties:
 HealthCheckIntervalSeconds: 10
 HealthCheckPath: "/readyz"
 HealthCheckPort: 6443
 HealthCheckProtocol: HTTPS
 HealthyThresholdCount: 2
 UnhealthyThresholdCount: 2
 Port: 6443
 Protocol: TCP
 TargetType: ip
 VpcId:
 Ref: VpcId
 TargetGroupAttributes:
 - Key: deregistration_delay.timeout_seconds
 Value: 60

InternalServiceInternalListener:
Type: AWS::ElasticLoadBalancingV2::Listener
Properties:
 DefaultActions:
 - Type: forward
 TargetGroupArn:
 Ref: InternalServiceTargetGroup
 LoadBalancerArn:
 Ref: IntApiElb
 Port: 22623
 Protocol: TCP

InternalServiceTargetGroup:
Type: AWS::ElasticLoadBalancingV2::TargetGroup
Properties:
 HealthCheckIntervalSeconds: 10
 HealthCheckPath: "/healthz"
 HealthCheckPort: 22623
 HealthCheckProtocol: HTTPS
 HealthyThresholdCount: 2
 UnhealthyThresholdCount: 2
 Port: 22623
 Protocol: TCP
 TargetType: ip
 VpcId:
 Ref: VpcId
 TargetGroupAttributes:
 - Key: deregistration_delay.timeout_seconds
 Value: 60

```

RegisterTargetLambdalamRole:
  Type: AWS::IAM::Role
  Properties:
    RoleName: !Join ["-", [!Ref InfrastructureName, "nlb", "lambda", "role"]]
    AssumeRolePolicyDocument:
      Version: "2012-10-17"
      Statement:
        - Effect: "Allow"
          Principal:
            Service:
              - "lambda.amazonaws.com"
          Action:
            - "sts:AssumeRole"
    Path: "/"
  Policies:
    - PolicyName: !Join ["-", [!Ref InfrastructureName, "master", "policy"]]
      PolicyDocument:
        Version: "2012-10-17"
        Statement:
          - Effect: "Allow"
            Action:
              [
                "elasticloadbalancing:RegisterTargets",
                "elasticloadbalancing:DeregisterTargets",
              ]
            Resource: !Ref InternalApiTargetGroup
        - Effect: "Allow"
          Action:
            [
              "elasticloadbalancing:RegisterTargets",
              "elasticloadbalancing:DeregisterTargets",
            ]
          Resource: !Ref InternalServiceTargetGroup
        - Effect: "Allow"
          Action:
            [
              "elasticloadbalancing:RegisterTargets",
              "elasticloadbalancing:DeregisterTargets",
            ]
          Resource: !Ref ExternalApiTargetGroup

```

```

RegisterNlbTargets:
  Type: "AWS::Lambda::Function"
  Properties:
    Handler: "index.handler"
    Role:
      Fn::GetAtt:
        - "RegisterTargetLambdalamRole"
        - "Arn"
  Code:
    ZipFile: |
      import json
      import boto3
      import cfnresponse
      def handler(event, context):

```

```

elb = boto3.client('elbv2')
if event['RequestType'] == 'Delete':
    elb.deregister_targets(TargetGroupArn=event['ResourceProperties']
['TargetArn'],Targets=[{'Id': event['ResourceProperties']['TargetIp']})
    elif event['RequestType'] == 'Create':
        elb.register_targets(TargetGroupArn=event['ResourceProperties']['TargetArn'],Targets=
[{'Id': event['ResourceProperties']['TargetIp']})
        responseData = {}
        cfnresponse.send(event, context, cfnresponse.SUCCESS, responseData,
event['ResourceProperties']['TargetArn']+event['ResourceProperties']['TargetIp'])
    Runtime: "python3.8"
    Timeout: 120

```

RegisterSubnetTagsLambdalamRole:

Type: AWS::IAM::Role

Properties:

RoleName: !Join ["-", [!Ref InfrastructureName, "subnet-tags-lambda-role"]]

AssumeRolePolicyDocument:

Version: "2012-10-17"

Statement:

- Effect: "Allow"

Principal:

Service:

- "lambda.amazonaws.com"

Action:

- "sts:AssumeRole"

Path: "/"

Policies:

- PolicyName: !Join ["-", [!Ref InfrastructureName, "subnet-tagging-policy"]]

PolicyDocument:

Version: "2012-10-17"

Statement:

- Effect: "Allow"

Action:

```

[
  "ec2:DeleteTags",
  "ec2:CreateTags"
]

```

Resource: "arn:aws:ec2:*:*:subnet/*"

- Effect: "Allow"

Action:

```

[
  "ec2:DescribeSubnets",
  "ec2:DescribeTags"
]

```

Resource: ""

RegisterSubnetTags:

Type: "AWS::Lambda::Function"

Properties:

Handler: "index.handler"

Role:

Fn::GetAtt:

- "RegisterSubnetTagsLambdalamRole"

- "Arn"

Code:

```

ZipFile: |
import json
import boto3
import cfnresponse
def handler(event, context):
    ec2_client = boto3.client('ec2')
    if event['RequestType'] == 'Delete':
        for subnet_id in event['ResourceProperties']['Subnets']:
            ec2_client.delete_tags(Resources=[subnet_id], Tags=[{'Key': 'kubernetes.io/cluster/' +
event['ResourceProperties']['InfrastructureName']}]);
    elif event['RequestType'] == 'Create':
        for subnet_id in event['ResourceProperties']['Subnets']:
            ec2_client.create_tags(Resources=[subnet_id], Tags=[{'Key': 'kubernetes.io/cluster/' +
event['ResourceProperties']['InfrastructureName'], 'Value': 'shared'}]);
    responseData = {}
    cfnresponse.send(event, context, cfnresponse.SUCCESS, responseData,
event['ResourceProperties']['InfrastructureName']+event['ResourceProperties']['Subnets'][0])
Runtime: "python3.8"
Timeout: 120

```

```

RegisterPublicSubnetTags:
Type: Custom::SubnetRegister
Properties:
ServiceToken: !GetAtt RegisterSubnetTags.Arn
InfrastructureName: !Ref InfrastructureName
Subnets: !Ref PublicSubnets

```

```

RegisterPrivateSubnetTags:
Type: Custom::SubnetRegister
Properties:
ServiceToken: !GetAtt RegisterSubnetTags.Arn
InfrastructureName: !Ref InfrastructureName
Subnets: !Ref PrivateSubnets

```

Outputs:

```

PrivateHostedZoneId:
Description: Hosted zone ID for the private DNS, which is required for private records.
Value: !Ref IntDns
ExternalApiLoadBalancerName:
Description: Full name of the external API load balancer.
Value: !GetAtt ExtApiElb.LoadBalancerFullName
InternalApiLoadBalancerName:
Description: Full name of the internal API load balancer.
Value: !GetAtt IntApiElb.LoadBalancerFullName
ApiServerDnsName:
Description: Full hostname of the API server, which is required for the Ignition config files.
Value: !Join [".", ["api-int", !Ref ClusterName, !Ref HostedZoneName]]
RegisterNlbIpTargetsLambda:
Description: Lambda ARN useful to help register or deregister IP targets for these load
balancers.
Value: !GetAtt RegisterNlbIpTargets.Arn
ExternalApiTargetGroupArn:
Description: ARN of the external API target group.
Value: !Ref ExternalApiTargetGroup
InternalApiTargetGroupArn:
Description: ARN of the internal API target group.

```


Value: !Ref InternalApiTargetGroup
 InternalServiceTargetGroupArn:
 Description: ARN of the internal service target group.
 Value: !Ref InternalServiceTargetGroup

IMPORTANT

If you are deploying your cluster to an AWS government or secret region, you must update the **InternalApiServerRecord** to use **CNAME** records. Records of type **ALIAS** are not supported for AWS government regions. For example:

Type: CNAME
 TTL: 10
 ResourceRecords:
 - !GetAtt IntApiElb.DNSName

Additional resources

- You can view details about the CloudFormation stacks that you create by navigating to the [AWS CloudFormation console](#).
- You can view details about your hosted zones by navigating to the [AWS Route 53 console](#).
- See [Listing public hosted zones](#) in the AWS documentation for more information about listing public hosted zones.

6.4.3.6. Creating security group and roles in AWS

You must create security groups and roles in Amazon Web Services (AWS) for your OpenShift Container Platform cluster to use.

You can use the provided CloudFormation template and a custom parameter file to create a stack of AWS resources. The stack represents the security groups and roles that your OpenShift Container Platform cluster requires.

NOTE

If you do not use the provided CloudFormation template to create your AWS infrastructure, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites


- You configured an AWS account.
- You added your AWS keys and region to your local AWS profile by running **aws configure**.
- You generated the Ignition config files for your cluster.
- You created and configured a VPC and associated subnets in AWS.

Procedure

1. Create a JSON file that contains the parameter values that the template requires:

```
[
  {
    "ParameterKey": "InfrastructureName", 1
    "ParameterValue": "mycluster-<random_string>" 2
  },
  {
    "ParameterKey": "VpcCidr", 3
    "ParameterValue": "10.0.0.0/16" 4
  },
  {
    "ParameterKey": "PrivateSubnets", 5
    "ParameterValue": "subnet-<random_string>" 6
  },
  {
    "ParameterKey": "VpcId", 7
    "ParameterValue": "vpc-<random_string>" 8
  }
]
```

- 1 The name for your cluster infrastructure that is encoded in your Ignition config files for the cluster.
 - 2 Specify the infrastructure name that you extracted from the Ignition config file metadata, which has the format **<cluster-name>-<random-string>**.
 - 3 The CIDR block for the VPC.
 - 4 Specify the CIDR block parameter that you used for the VPC that you defined in the form **x.x.x.x/16-24**.
 - 5 The private subnets that you created for your VPC.
 - 6 Specify the **PrivateSubnetIds** value from the output of the CloudFormation template for the VPC.
 - 7 The VPC that you created for the cluster.
 - 8 Specify the **VpcId** value from the output of the CloudFormation template for the VPC.
2. Copy the template from the **CloudFormation template for security objects** section of this topic and save it as a YAML file on your computer. This template describes the security groups and roles that your cluster requires.
 3. Launch the CloudFormation template to create a stack of AWS resources that represent the security groups and roles:



IMPORTANT

You must enter the command on a single line.



```
$ aws cloudformation create-stack --stack-name <name> 1
  --template-body file://<template>.yaml 2
  --parameters file://<parameters>.json 3
  --capabilities CAPABILITY_NAMED_IAM 4
```

- 1 **<name>** is the name for the CloudFormation stack, such as **cluster-sec**. You need the name of this stack if you remove the cluster.
- 2 **<template>** is the relative path to and name of the CloudFormation template YAML file that you saved.
- 3 **<parameters>** is the relative path to and name of the CloudFormation parameters JSON file.
- 4 You must explicitly declare the **CAPABILITY_NAMED_IAM** capability because the provided template creates some **AWS::IAM::Role** and **AWS::IAM::InstanceProfile** resources.

Example output

```
arn:aws:cloudformation:us-east-1:269333783861:stack/cluster-sec/03bd4210-2ed7-11eb-6d7a-13fc0b61e9db
```

4. Confirm that the template components exist:

```
$ aws cloudformation describe-stacks --stack-name <name>
```

After the **StackStatus** displays **CREATE_COMPLETE**, the output displays values for the following parameters. You must provide these parameter values to the other CloudFormation templates that you run to create your cluster:

MasterSecurityGroupID	Master Security Group ID
WorkerSecurityGroupID	Worker Security Group ID
MasterInstanceProfile	Master IAM Instance Profile
WorkerInstanceProfile	Worker IAM Instance Profile

6.4.3.6.1. CloudFormation template for security objects

You can use the following CloudFormation template to deploy the security objects that you need for your OpenShift Container Platform cluster.

Example 6.74. CloudFormation template for security objects

AWSTemplateFormatVersion: 2010-09-09

Description: Template for OpenShift Cluster Security Elements (Security Groups & IAM)

Parameters:

InfrastructureName:

AllowedPattern: ^([a-zA-Z][a-zA-Z0-9\-_]{0,26})\$

MaxLength: 27

MinLength: 1

ConstraintDescription: Infrastructure name must be alphanumeric, start with a letter, and have a maximum of 27 characters.

Description: A short, unique cluster ID used to tag cloud resources and identify items owned or used by the cluster.

Type: String

VpcCidr:

AllowedPattern: ^((([0-9]|[1-9][0-9]|1[0-9]{2}|2[0-4][0-9]|25[0-5])\.)\.)\{3\}([0-9]|[1-9][0-9]|1[0-9]{2}|2[0-4][0-9]|25[0-5])\)(\^(1[6-9]|2[0-4]))\$

ConstraintDescription: CIDR block parameter must be in the form x.x.x.x/16-24.

Default: 10.0.0/16

Description: CIDR block for VPC.

Type: String

VpcId:

Description: The VPC-scoped resources will belong to this VPC.

Type: AWS::EC2::VPC::Id

PrivateSubnets:

Description: The internal subnets.

Type: List<AWS::EC2::Subnet::Id>

Metadata:

AWS::CloudFormation::Interface:

ParameterGroups:

- Label:

default: "Cluster Information"

Parameters:

- InfrastructureName

- Label:

default: "Network Configuration"

Parameters:

- VpcId

- VpcCidr

- PrivateSubnets

ParameterLabels:

InfrastructureName:

default: "Infrastructure Name"

VpcId:

default: "VPC ID"

VpcCidr:

default: "VPC CIDR"

PrivateSubnets:

default: "Private Subnets"

Resources:

MasterSecurityGroup:

Type: AWS::EC2::SecurityGroup

Properties:

GroupDescription: Cluster Master Security Group

SecurityGroupIngress:

- IpProtocol: icmp
 - FromPort: 0
 - ToPort: 0
 - CidrIp: !Ref VpcCidr
 - IpProtocol: tcp
 - FromPort: 22
 - ToPort: 22
 - CidrIp: !Ref VpcCidr
 - IpProtocol: tcp
 - ToPort: 6443
 - FromPort: 6443
 - CidrIp: !Ref VpcCidr
 - IpProtocol: tcp
 - FromPort: 22623
 - ToPort: 22623
 - CidrIp: !Ref VpcCidr
- VpcId: !Ref VpcId

WorkerSecurityGroup:

Type: AWS::EC2::SecurityGroup

Properties:

GroupDescription: Cluster Worker Security Group

SecurityGroupIngress:

- IpProtocol: icmp
 - FromPort: 0
 - ToPort: 0
 - CidrIp: !Ref VpcCidr
 - IpProtocol: tcp
 - FromPort: 22
 - ToPort: 22
 - CidrIp: !Ref VpcCidr
- VpcId: !Ref VpcId

MasterIngressEtcD:

Type: AWS::EC2::SecurityGroupIngress

Properties:

GroupId: !GetAtt MasterSecurityGroup.GroupId
 SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId
 Description: etcd
 FromPort: 2379
 ToPort: 2380
 IpProtocol: tcp

MasterIngressVxlan:

Type: AWS::EC2::SecurityGroupIngress

Properties:

GroupId: !GetAtt MasterSecurityGroup.GroupId
 SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId
 Description: Vxlan packets
 FromPort: 4789
 ToPort: 4789
 IpProtocol: udp

MasterIngressWorkerVxlan:

Type: AWS::EC2::SecurityGroupIngress
Properties:
 GroupId: !GetAtt MasterSecurityGroup.GroupId
 SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId
 Description: Vxlan packets
 FromPort: 4789
 ToPort: 4789
 IpProtocol: udp

MasterIngressGeneve:
Type: AWS::EC2::SecurityGroupIngress
Properties:
 GroupId: !GetAtt MasterSecurityGroup.GroupId
 SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId
 Description: Geneve packets
 FromPort: 6081
 ToPort: 6081
 IpProtocol: udp

MasterIngressWorkerGeneve:
Type: AWS::EC2::SecurityGroupIngress
Properties:
 GroupId: !GetAtt MasterSecurityGroup.GroupId
 SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId
 Description: Geneve packets
 FromPort: 6081
 ToPort: 6081
 IpProtocol: udp

MasterIngressIpsecIke:
Type: AWS::EC2::SecurityGroupIngress
Properties:
 GroupId: !GetAtt MasterSecurityGroup.GroupId
 SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId
 Description: IPsec IKE packets
 FromPort: 500
 ToPort: 500
 IpProtocol: udp

MasterIngressIpsecNat:
Type: AWS::EC2::SecurityGroupIngress
Properties:
 GroupId: !GetAtt MasterSecurityGroup.GroupId
 SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId
 Description: IPsec NAT-T packets
 FromPort: 4500
 ToPort: 4500
 IpProtocol: udp

MasterIngressIpsecEsp:
Type: AWS::EC2::SecurityGroupIngress
Properties:
 GroupId: !GetAtt MasterSecurityGroup.GroupId
 SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId
 Description: IPsec ESP packets
 IpProtocol: 50

MasterIngressWorkerIpsecIke:

Type: AWS::EC2::SecurityGroupIngress

Properties:

GroupId: !GetAtt MasterSecurityGroup.GroupId

SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId

Description: IPsec IKE packets

FromPort: 500

ToPort: 500

IpProtocol: udp

MasterIngressWorkerIpsecNat:

Type: AWS::EC2::SecurityGroupIngress

Properties:

GroupId: !GetAtt MasterSecurityGroup.GroupId

SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId

Description: IPsec NAT-T packets

FromPort: 4500

ToPort: 4500

IpProtocol: udp

MasterIngressWorkerIpsecEsp:

Type: AWS::EC2::SecurityGroupIngress

Properties:

GroupId: !GetAtt MasterSecurityGroup.GroupId

SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId

Description: IPsec ESP packets

IpProtocol: 50

MasterIngressInternal:

Type: AWS::EC2::SecurityGroupIngress

Properties:

GroupId: !GetAtt MasterSecurityGroup.GroupId

SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId

Description: Internal cluster communication

FromPort: 9000

ToPort: 9999

IpProtocol: tcp

MasterIngressWorkerInternal:

Type: AWS::EC2::SecurityGroupIngress

Properties:

GroupId: !GetAtt MasterSecurityGroup.GroupId

SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId

Description: Internal cluster communication

FromPort: 9000

ToPort: 9999

IpProtocol: tcp

MasterIngressInternalUDP:

Type: AWS::EC2::SecurityGroupIngress

Properties:

GroupId: !GetAtt MasterSecurityGroup.GroupId

SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId

Description: Internal cluster communication

FromPort: 9000

ToPort: 9999
IpProtocol: udp

MasterIngressWorkerInternalUDP:

Type: AWS::EC2::SecurityGroupIngress
Properties:
GroupId: !GetAtt MasterSecurityGroup.GroupId
SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId
Description: Internal cluster communication
FromPort: 9000
ToPort: 9999
IpProtocol: udp

MasterIngressKube:

Type: AWS::EC2::SecurityGroupIngress
Properties:
GroupId: !GetAtt MasterSecurityGroup.GroupId
SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId
Description: Kubernetes kubelet, scheduler and controller manager
FromPort: 10250
ToPort: 10259
IpProtocol: tcp

MasterIngressWorkerKube:

Type: AWS::EC2::SecurityGroupIngress
Properties:
GroupId: !GetAtt MasterSecurityGroup.GroupId
SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId
Description: Kubernetes kubelet, scheduler and controller manager
FromPort: 10250
ToPort: 10259
IpProtocol: tcp

MasterIngressIngressServices:

Type: AWS::EC2::SecurityGroupIngress
Properties:
GroupId: !GetAtt MasterSecurityGroup.GroupId
SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId
Description: Kubernetes ingress services
FromPort: 30000
ToPort: 32767
IpProtocol: tcp

MasterIngressWorkerIngressServices:

Type: AWS::EC2::SecurityGroupIngress
Properties:
GroupId: !GetAtt MasterSecurityGroup.GroupId
SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId
Description: Kubernetes ingress services
FromPort: 30000
ToPort: 32767
IpProtocol: tcp

MasterIngressIngressServicesUDP:

Type: AWS::EC2::SecurityGroupIngress
Properties:

GroupId: !GetAtt MasterSecurityGroup.GroupId
SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId
Description: Kubernetes ingress services
FromPort: 30000
ToPort: 32767
IpProtocol: udp

MasterIngressWorkerIngressServicesUDP:
Type: AWS::EC2::SecurityGroupIngress
Properties:
GroupId: !GetAtt MasterSecurityGroup.GroupId
SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId
Description: Kubernetes ingress services
FromPort: 30000
ToPort: 32767
IpProtocol: udp

WorkerIngressVxlan:
Type: AWS::EC2::SecurityGroupIngress
Properties:
GroupId: !GetAtt WorkerSecurityGroup.GroupId
SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId
Description: Vxlan packets
FromPort: 4789
ToPort: 4789
IpProtocol: udp

WorkerIngressMasterVxlan:
Type: AWS::EC2::SecurityGroupIngress
Properties:
GroupId: !GetAtt WorkerSecurityGroup.GroupId
SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId
Description: Vxlan packets
FromPort: 4789
ToPort: 4789
IpProtocol: udp

WorkerIngressGeneve:
Type: AWS::EC2::SecurityGroupIngress
Properties:
GroupId: !GetAtt WorkerSecurityGroup.GroupId
SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId
Description: Geneve packets
FromPort: 6081
ToPort: 6081
IpProtocol: udp

WorkerIngressMasterGeneve:
Type: AWS::EC2::SecurityGroupIngress
Properties:
GroupId: !GetAtt WorkerSecurityGroup.GroupId
SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId
Description: Geneve packets
FromPort: 6081
ToPort: 6081
IpProtocol: udp

WorkerIngressIpsecIke:

Type: AWS::EC2::SecurityGroupIngress

Properties:

GroupId: !GetAtt WorkerSecurityGroup.GroupId

SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId

Description: IPsec IKE packets

FromPort: 500

ToPort: 500

IpProtocol: udp

WorkerIngressIpsecNat:

Type: AWS::EC2::SecurityGroupIngress

Properties:

GroupId: !GetAtt WorkerSecurityGroup.GroupId

SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId

Description: IPsec NAT-T packets

FromPort: 4500

ToPort: 4500

IpProtocol: udp

WorkerIngressIpsecEsp:

Type: AWS::EC2::SecurityGroupIngress

Properties:

GroupId: !GetAtt WorkerSecurityGroup.GroupId

SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId

Description: IPsec ESP packets

IpProtocol: 50

WorkerIngressMasterIpsecIke:

Type: AWS::EC2::SecurityGroupIngress

Properties:

GroupId: !GetAtt WorkerSecurityGroup.GroupId

SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId

Description: IPsec IKE packets

FromPort: 500

ToPort: 500

IpProtocol: udp

WorkerIngressMasterIpsecNat:

Type: AWS::EC2::SecurityGroupIngress

Properties:

GroupId: !GetAtt WorkerSecurityGroup.GroupId

SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId

Description: IPsec NAT-T packets

FromPort: 4500

ToPort: 4500

IpProtocol: udp

WorkerIngressMasterIpsecEsp:

Type: AWS::EC2::SecurityGroupIngress

Properties:

GroupId: !GetAtt WorkerSecurityGroup.GroupId

SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId

Description: IPsec ESP packets

IpProtocol: 50

WorkerIngressInternal:

Type: AWS::EC2::SecurityGroupIngress

Properties:

GroupId: !GetAtt WorkerSecurityGroup.GroupId

SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId

Description: Internal cluster communication

FromPort: 9000

ToPort: 9999

IpProtocol: tcp

WorkerIngressMasterInternal:

Type: AWS::EC2::SecurityGroupIngress

Properties:

GroupId: !GetAtt WorkerSecurityGroup.GroupId

SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId

Description: Internal cluster communication

FromPort: 9000

ToPort: 9999

IpProtocol: tcp

WorkerIngressInternalUDP:

Type: AWS::EC2::SecurityGroupIngress

Properties:

GroupId: !GetAtt WorkerSecurityGroup.GroupId

SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId

Description: Internal cluster communication

FromPort: 9000

ToPort: 9999

IpProtocol: udp

WorkerIngressMasterInternalUDP:

Type: AWS::EC2::SecurityGroupIngress

Properties:

GroupId: !GetAtt WorkerSecurityGroup.GroupId

SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId

Description: Internal cluster communication

FromPort: 9000

ToPort: 9999

IpProtocol: udp

WorkerIngressKube:

Type: AWS::EC2::SecurityGroupIngress

Properties:

GroupId: !GetAtt WorkerSecurityGroup.GroupId

SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId

Description: Kubernetes secure kubelet port

FromPort: 10250

ToPort: 10250

IpProtocol: tcp

WorkerIngressWorkerKube:

Type: AWS::EC2::SecurityGroupIngress

Properties:

GroupId: !GetAtt WorkerSecurityGroup.GroupId

SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId

Description: Internal Kubernetes communication

FromPort: 10250

ToPort: 10250

IpProtocol: tcp

WorkerIngressIngressServices:

Type: AWS::EC2::SecurityGroupIngress

Properties:

GroupId: !GetAtt WorkerSecurityGroup.GroupId

SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId

Description: Kubernetes ingress services

FromPort: 30000

ToPort: 32767

IpProtocol: tcp

WorkerIngressMasterIngressServices:

Type: AWS::EC2::SecurityGroupIngress

Properties:

GroupId: !GetAtt WorkerSecurityGroup.GroupId

SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId

Description: Kubernetes ingress services

FromPort: 30000

ToPort: 32767

IpProtocol: tcp

WorkerIngressIngressServicesUDP:

Type: AWS::EC2::SecurityGroupIngress

Properties:

GroupId: !GetAtt WorkerSecurityGroup.GroupId

SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId

Description: Kubernetes ingress services

FromPort: 30000

ToPort: 32767

IpProtocol: udp

WorkerIngressMasterIngressServicesUDP:

Type: AWS::EC2::SecurityGroupIngress

Properties:

GroupId: !GetAtt WorkerSecurityGroup.GroupId

SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId

Description: Kubernetes ingress services

FromPort: 30000

ToPort: 32767

IpProtocol: udp

MasterIamRole:

Type: AWS::IAM::Role

Properties:

AssumeRolePolicyDocument:

Version: "2012-10-17"

Statement:

- Effect: "Allow"

Principal:

Service:

- "ec2.amazonaws.com"

Action:

- "sts:AssumeRole"

Policies:

- PolicyName: !Join ["-", [!Ref InfrastructureName, "master", "policy"]]

PolicyDocument:

Version: "2012-10-17"

Statement:

- Effect: "Allow"

Action:

- "ec2:AttachVolume"

- "ec2:AuthorizeSecurityGroupIngress"

- "ec2:CreateSecurityGroup"

- "ec2:CreateTags"

- "ec2:CreateVolume"

- "ec2>DeleteSecurityGroup"

- "ec2>DeleteVolume"

- "ec2:Describe*"

- "ec2:DetachVolume"

- "ec2:ModifyInstanceAttribute"

- "ec2:ModifyVolume"

- "ec2:RevokeSecurityGroupIngress"

- "elasticloadbalancing:AddTags"

- "elasticloadbalancing:AttachLoadBalancerToSubnets"

- "elasticloadbalancing:ApplySecurityGroupsToLoadBalancer"

- "elasticloadbalancing:CreateListener"

- "elasticloadbalancing:CreateLoadBalancer"

- "elasticloadbalancing:CreateLoadBalancerPolicy"

- "elasticloadbalancing:CreateLoadBalancerListeners"

- "elasticloadbalancing:CreateTargetGroup"

- "elasticloadbalancing:ConfigureHealthCheck"

- "elasticloadbalancing>DeleteListener"

- "elasticloadbalancing>DeleteLoadBalancer"

- "elasticloadbalancing>DeleteLoadBalancerListeners"

- "elasticloadbalancing>DeleteTargetGroup"

- "elasticloadbalancing:DeregisterInstancesFromLoadBalancer"

- "elasticloadbalancing:DeregisterTargets"

- "elasticloadbalancing:Describe*"

- "elasticloadbalancing:DetachLoadBalancerFromSubnets"

- "elasticloadbalancing:ModifyListener"

- "elasticloadbalancing:ModifyLoadBalancerAttributes"

- "elasticloadbalancing:ModifyTargetGroup"

- "elasticloadbalancing:ModifyTargetGroupAttributes"

- "elasticloadbalancing:RegisterInstancesWithLoadBalancer"

- "elasticloadbalancing:RegisterTargets"

- "elasticloadbalancing:SetLoadBalancerPoliciesForBackendServer"

- "elasticloadbalancing:SetLoadBalancerPoliciesOfListener"

- "kms:DescribeKey"

Resource: ""

MasterInstanceProfile:

Type: "AWS::IAM::InstanceProfile"

Properties:

Roles:

- Ref: "MasterIamRole"

WorkerIamRole:

Type: AWS::IAM::Role

```

Properties:
  AssumeRolePolicyDocument:
    Version: "2012-10-17"
    Statement:
      - Effect: "Allow"
        Principal:
          Service:
            - "ec2.amazonaws.com"
        Action:
          - "sts:AssumeRole"
  Policies:
    - PolicyName: !Join ["-", [!Ref InfrastructureName, "worker", "policy"]]
      PolicyDocument:
        Version: "2012-10-17"
        Statement:
          - Effect: "Allow"
            Action:
              - "ec2:DescribeInstances"
              - "ec2:DescribeRegions"
            Resource: "*"

WorkerInstanceProfile:
  Type: "AWS::IAM::InstanceProfile"
  Properties:
    Roles:
      - Ref: "WorkerIamRole"

Outputs:
  MasterSecurityGroupId:
    Description: Master Security Group ID
    Value: !GetAtt MasterSecurityGroup.GroupId

  WorkerSecurityGroupId:
    Description: Worker Security Group ID
    Value: !GetAtt WorkerSecurityGroup.GroupId

  MasterInstanceProfile:
    Description: Master IAM Instance Profile
    Value: !Ref MasterInstanceProfile

  WorkerInstanceProfile:
    Description: Worker IAM Instance Profile
    Value: !Ref WorkerInstanceProfile

```

Additional resources

- You can view details about the CloudFormation stacks that you create by navigating to the [AWS CloudFormation console](#).

6.4.3.7. Accessing RHCOS AMIs with stream metadata

In OpenShift Container Platform, *stream metadata* provides standardized metadata about RHCOS in the JSON format and injects the metadata into the cluster. Stream metadata is a stable format that supports multiple architectures and is intended to be self-documenting for maintaining automation.

You can use the **coreos print-stream-json** sub-command of **openshift-install** to access information about the boot images in the stream metadata format. This command provides a method for printing stream metadata in a scriptable, machine-readable format.

For user-provisioned installations, the **openshift-install** binary contains references to the version of RHCOS boot images that are tested for use with OpenShift Container Platform, such as the AWS AMI.

Procedure

To parse the stream metadata, use one of the following methods:

- From a Go program, use the official **stream-metadata-go** library at <https://github.com/coreos/stream-metadata-go>. You can also view example code in the library.
- From another programming language, such as Python or Ruby, use the JSON library of your preferred programming language.
- From a command-line utility that handles JSON data, such as **jq**:
 - Print the current **x86_64** or **aarch64** AMI for an AWS region, such as **us-west-1**:

For x86_64

```
$ openshift-install coreos print-stream-json | jq -r
'.architectures.x86_64.images.aws.regions["us-west-1"].image'
```

Example output

```
ami-0d3e625f84626bbda
```

For aarch64

```
$ openshift-install coreos print-stream-json | jq -r
'.architectures.aarch64.images.aws.regions["us-west-1"].image'
```


Example output

```
ami-0af1d3b7fa5be2131
```

The output of this command is the AWS AMI ID for your designated architecture and the **us-west-1** region. The AMI must belong to the same region as the cluster.

6.4.3.8. RHCOS AMIs for the AWS infrastructure

Red Hat provides Red Hat Enterprise Linux CoreOS (RHCOS) AMIs that are valid for the various AWS regions and instance architectures that you can manually specify for your OpenShift Container Platform nodes.



NOTE

By importing your own AMI, you can also install to regions that do not have a published RHCOS AMI.

Table 6.23. x86_64 RHCOS AMIs

AWS zone	AWS AMI
af-south-1	ami-018de6b1be470b7e6
ap-east-1	ami-092f09a4c8aacf56a
ap-northeast-1	ami-001c9fc85b77505f4
ap-northeast-2	ami-0a5d7f1966d88fba3
ap-northeast-3	ami-085a2726ceb4f5eeb
ap-south-1	ami-03f2c19071fb6eab3
ap-south-2	ami-0c82522890979d94b
ap-southeast-1	ami-06e5b9d16ead6c55c
ap-southeast-2	ami-0ccd429129fbf6bc0
ap-southeast-3	ami-096f9b4d41116d95c
ap-southeast-4	ami-0b511ab55f9f7d052
ca-central-1	ami-0e357287c651be1f2
ca-west-1	ami-0b1692e5776740901
eu-central-1	ami-08b13fb388ba5610d
eu-central-2	ami-04777298b55045ea9
eu-north-1	ami-05421993a4ee567b9
eu-south-1	ami-00959341869d34746
eu-south-2	ami-0a745b610522a87cc
eu-west-1	ami-0e48f85ec57bd7baa
eu-west-2	ami-0c015b1387acddc5e
eu-west-3	ami-01e466af77a95b93d
il-central-1	ami-0a1b706793b54d087

AWS zone	AWS AMI
me-central-1	ami-09d58e8b2099ae5bc
me-south-1	ami-0f9c259bc4346599c
sa-east-1	ami-06277517d966881a3
us-east-1	ami-05483066c3caaccf5
us-east-2	ami-0c704ce516493a479
us-gov-east-1	ami-0695b2cbdd1ec4d48
us-gov-west-1	ami-004404a0eaa427b39
us-west-1	ami-0aaa56637063be772
us-west-2	ami-0870c7a3d5ef4d2b3

Table 6.24. aarch64 RHCOS AMIs

AWS zone	AWS AMI
af-south-1	ami-0d96d447d3df14f4e
ap-east-1	ami-010236402dfba1717
ap-northeast-1	ami-08feeb6d9068bb12e
ap-northeast-2	ami-0772082c119147205
ap-northeast-3	ami-05bcbb13493d0516f
ap-south-1	ami-0034a539e8adcb817
ap-south-2	ami-0fc56b7c53c38ff70
ap-southeast-1	ami-0b50a0e92a58f3ad8
ap-southeast-2	ami-0697a9174ae206182
ap-southeast-3	ami-05ae309319a7ad504
ap-southeast-4	ami-00500414843baa657

AWS zone	AWS AMI
ca-central-1	ami-05e76bf99d3b0d4c8
ca-west-1	ami-099fc953c221ec590
eu-central-1	ami-0d2e2698884020b71
eu-central-2	ami-0a96ba21ddee01719
eu-north-1	ami-0ab8a1070d0b1d9a5
eu-south-1	ami-0d0465299a8b196c1
eu-south-2	ami-07d5aa3c6d468c738
eu-west-1	ami-08e7d209f26c09c80
eu-west-2	ami-0c8336d4e2c1f13cb
eu-west-3	ami-085d804b98acf0648
il-central-1	ami-02ce030e36aeb7643
me-central-1	ami-0dea3ac466040b77f
me-south-1	ami-02ef2a46887b9786d
sa-east-1	ami-0d19817d31b2399f9
us-east-1	ami-04859c888566a2c2f
us-east-2	ami-02f7e4a5dc66361bb
us-gov-east-1	ami-067ef782980ea96ad
us-gov-west-1	ami-0ce67540c1bb2319d
us-west-1	ami-0a09c5d2f287718a3
us-west-2	ami-06b94e64e595f6cee

6.4.3.8.1. AWS regions without a published RHCOS AMI

You can deploy an OpenShift Container Platform cluster to Amazon Web Services (AWS) regions without native support for a Red Hat Enterprise Linux CoreOS (RHCOS) Amazon Machine Image (AMI) or the AWS software development kit (SDK). If a published AMI is not available for an AWS region, you

can upload a custom AMI prior to installing the cluster.

If you are deploying to a region not supported by the AWS SDK and you do not specify a custom AMI, the installation program copies the **us-east-1** AMI to the user account automatically. Then the installation program creates the control plane machines with encrypted EBS volumes using the default or user-specified Key Management Service (KMS) key. This allows the AMI to follow the same process workflow as published RHCOS AMIs.

A region without native support for an RHCOS AMI is not available to select from the terminal during cluster creation because it is not published. However, you can install to this region by configuring the custom AMI in the **install-config.yaml** file.

6.4.3.8.2. Uploading a custom RHCOS AMI in AWS

If you are deploying to a custom Amazon Web Services (AWS) region, you must upload a custom Red Hat Enterprise Linux CoreOS (RHCOS) Amazon Machine Image (AMI) that belongs to that region.

Prerequisites

- You configured an AWS account.
- You created an Amazon S3 bucket with the required IAM [service role](#).
- You uploaded your RHCOS VMDK file to Amazon S3. The RHCOS VMDK file must be the highest version that is less than or equal to the OpenShift Container Platform version you are installing.
- You downloaded the AWS CLI and installed it on your computer. See [Install the AWS CLI Using the Bundled Installer](#).

Procedure

1. Export your AWS profile as an environment variable:

```
$ export AWS_PROFILE=<aws_profile> 1
```

2. Export the region to associate with your custom AMI as an environment variable:

```
$ export AWS_DEFAULT_REGION=<aws_region> 1
```

3. Export the version of RHCOS you uploaded to Amazon S3 as an environment variable:

```
$ export RHCOS_VERSION=<version> 1
```

1 1 1 The RHCOS VMDK version, like **4.16.0**.

4. Export the Amazon S3 bucket name as an environment variable:

```
$ export VMIMPORT_BUCKET_NAME=<s3_bucket_name>
```

5. Create the **containers.json** file and define your RHCOS VMDK file:

```
$ cat <<EOF > containers.json
```

```
{
  "Description": "rhcos-${RHCOS_VERSION}-x86_64-aws.x86_64",
  "Format": "vmdk",
  "UserBucket": {
    "S3Bucket": "${VMIMPORT_BUCKET_NAME}",
    "S3Key": "rhcos-${RHCOS_VERSION}-x86_64-aws.x86_64.vmdk"
  }
}
EOF
```

6. Import the RHCOS disk as an Amazon EBS snapshot:

```
$ aws ec2 import-snapshot --region ${AWS_DEFAULT_REGION} \
  --description "<description>" ❶
  --disk-container "file://<file_path>/containers.json" ❷
```

- ❶ The description of your RHCOS disk being imported, like **rhcos-\${RHCOS_VERSION}-x86_64-aws.x86_64**.
- ❷ The file path to the JSON file describing your RHCOS disk. The JSON file should contain your Amazon S3 bucket name and key.

7. Check the status of the image import:

```
$ watch -n 5 aws ec2 describe-import-snapshot-tasks --region ${AWS_DEFAULT_REGION}
```

Example output

```
{
  "ImportSnapshotTasks": [
    {
      "Description": "rhcos-4.7.0-x86_64-aws.x86_64",
      "ImportTaskId": "import-snap-fh6i8uil",
      "SnapshotTaskDetail": {
        "Description": "rhcos-4.7.0-x86_64-aws.x86_64",
        "DiskImageSize": 819056640.0,
        "Format": "VMDK",
        "SnapshotId": "snap-06331325870076318",
        "Status": "completed",
        "UserBucket": {
          "S3Bucket": "external-images",
          "S3Key": "rhcos-4.7.0-x86_64-aws.x86_64.vmdk"
        }
      }
    }
  ]
}
```

Copy the **SnapshotId** to register the image.

8. Create a custom RHCOS AMI from the RHCOS snapshot:

```
$ aws ec2 register-image \
```

```

--region ${AWS_DEFAULT_REGION} \
--architecture x86_64 \ 1
--description "rhcos-${RHCOS_VERSION}-x86_64-aws.x86_64" \ 2
--ena-support \
--name "rhcos-${RHCOS_VERSION}-x86_64-aws.x86_64" \ 3
--virtualization-type hvm \
--root-device-name '/dev/xvda' \
--block-device-mappings 'DeviceName=/dev/xvda,Ebs=
{DeleteOnTermination=true,SnapshotId=<snapshot_ID>}' 4

```

- 1** The RHCOS VMDK architecture type, like **x86_64**, **aarch64**, **s390x**, or **ppc64le**.
- 2** The **Description** from the imported snapshot.
- 3** The name of the RHCOS AMI.
- 4** The **SnapshotID** from the imported snapshot.

To learn more about these APIs, see the AWS documentation for [importing snapshots](#) and [creating EBS-backed AMIs](#).

6.4.3.9. Creating the bootstrap node in AWS

You must create the bootstrap node in Amazon Web Services (AWS) to use during OpenShift Container Platform cluster initialization. You do this by:

- Providing a location to serve the **bootstrap.ign** Ignition config file to your cluster. This file is located in your installation directory. The provided CloudFormation Template assumes that the Ignition config files for your cluster are served from an S3 bucket. If you choose to serve the files from another location, you must modify the templates.
- Using the provided CloudFormation template and a custom parameter file to create a stack of AWS resources. The stack represents the bootstrap node that your OpenShift Container Platform installation requires.



NOTE

If you do not use the provided CloudFormation template to create your bootstrap node, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- You configured an AWS account.
- You added your AWS keys and region to your local AWS profile by running **aws configure**.
- You generated the Ignition config files for your cluster.
- You created and configured a VPC and associated subnets in AWS.
- You created and configured DNS, load balancers, and listeners in AWS.

- You created the security groups and roles required for your cluster in AWS.

Procedure

- Create the bucket by running the following command:

```
$ aws s3 mb s3://<cluster-name>-infra 1
```

- 1** `<cluster-name>-infra` is the bucket name. When creating the `install-config.yaml` file, replace `<cluster-name>` with the name specified for the cluster.

You must use a presigned URL for your S3 bucket, instead of the `s3://` schema, if you are:

- Deploying to a region that has endpoints that differ from the AWS SDK.
- Deploying a proxy.
- Providing your own custom endpoints.

- Upload the `bootstrap.ign` Ignition config file to the bucket by running the following command:

```
$ aws s3 cp <installation_directory>/bootstrap.ign s3://<cluster-name>-infra/bootstrap.ign 1
```

- 1** For `<installation_directory>`, specify the path to the directory that you stored the installation files in.

- Verify that the file uploaded by running the following command:

```
$ aws s3 ls s3://<cluster-name>-infra/
```

Example output

```
2019-04-03 16:15:16 314878 bootstrap.ign
```



NOTE

The bootstrap Ignition config file does contain secrets, like X.509 keys. The following steps provide basic security for the S3 bucket. To provide additional security, you can enable an S3 bucket policy to allow only certain users, such as the OpenShift IAM user, to access objects that the bucket contains. You can avoid S3 entirely and serve your bootstrap Ignition config file from any address that the bootstrap machine can reach.

- Create a JSON file that contains the parameter values that the template requires:

```
[
  {
    "ParameterKey": "InfrastructureName", 1
    "ParameterValue": "mycluster-<random_string>" 2
  },
  {
```

```

    "ParameterKey": "RhcosAmi", 3
    "ParameterValue": "ami-<random_string>" 4
  },
  {
    "ParameterKey": "AllowedBootstrapSshCidr", 5
    "ParameterValue": "0.0.0.0/0" 6
  },
  {
    "ParameterKey": "PublicSubnet", 7
    "ParameterValue": "subnet-<random_string>" 8
  },
  {
    "ParameterKey": "MasterSecurityGroupId", 9
    "ParameterValue": "sg-<random_string>" 10
  },
  {
    "ParameterKey": "VpcId", 11
    "ParameterValue": "vpc-<random_string>" 12
  },
  {
    "ParameterKey": "BootstrapIgnitionLocation", 13
    "ParameterValue": "s3://<bucket_name>/bootstrap.ign" 14
  },
  {
    "ParameterKey": "AutoRegisterELB", 15
    "ParameterValue": "yes" 16
  },
  {
    "ParameterKey": "RegisterNlbPTargetsLambdaArn", 17
    "ParameterValue": "arn:aws:lambda:<aws_region>:<account_number>:function:
<dns_stack_name>-RegisterNlbPTargets-<random_string>" 18
  },
  {
    "ParameterKey": "ExternalApiTargetGroupArn", 19
    "ParameterValue": "arn:aws:elasticloadbalancing:<aws_region>:
<account_number>:targetgroup/<dns_stack_name>-Exter-<random_string>" 20
  },
  {
    "ParameterKey": "InternalApiTargetGroupArn", 21
    "ParameterValue": "arn:aws:elasticloadbalancing:<aws_region>:
<account_number>:targetgroup/<dns_stack_name>-Inter-<random_string>" 22
  },
  {
    "ParameterKey": "InternalServiceTargetGroupArn", 23
    "ParameterValue": "arn:aws:elasticloadbalancing:<aws_region>:
<account_number>:targetgroup/<dns_stack_name>-Inter-<random_string>" 24
  }
]


```

- 1 The name for your cluster infrastructure that is encoded in your Ignition config files for the cluster.
- 2 Specify the infrastructure name that you extracted from the Ignition config file metadata, which has the format `<cluster_name>-<random_string>`.

which has the format **<cluster-name>-<random-string>**.

- 3 Current Red Hat Enterprise Linux CoreOS (RHCOS) AMI to use for the bootstrap node based on your selected architecture.
- 4 Specify a valid **AWS::EC2::Image::Id** value.
- 5 CIDR block to allow SSH access to the bootstrap node.
- 6 Specify a CIDR block in the format **x.x.x.x/16-24**.
- 7 The public subnet that is associated with your VPC to launch the bootstrap node into.
- 8 Specify the **PublicSubnetIds** value from the output of the CloudFormation template for the VPC.
- 9 The master security group ID (for registering temporary rules)
- 10 Specify the **MasterSecurityGroupId** value from the output of the CloudFormation template for the security group and roles.
- 11 The VPC created resources will belong to.
- 12 Specify the **VpcId** value from the output of the CloudFormation template for the VPC.
- 13 Location to fetch bootstrap Ignition config file from.
- 14 Specify the S3 bucket and file name in the form **s3://<bucket_name>/bootstrap.ign**.
- 15 Whether or not to register a network load balancer (NLB).
- 16 Specify **yes** or **no**. If you specify **yes**, you must provide a Lambda Amazon Resource Name (ARN) value.
- 17 The ARN for NLB IP target registration lambda group.
- 18 Specify the **RegisterNlbTargetsLambda** value from the output of the CloudFormation template for DNS and load balancing. Use **arn:aws-us-gov** if deploying the cluster to an AWS GovCloud region.
- 19 The ARN for external API load balancer target group.
- 20 Specify the **ExternalApiTargetGroupArn** value from the output of the CloudFormation template for DNS and load balancing. Use **arn:aws-us-gov** if deploying the cluster to an AWS GovCloud region.
- 21 The ARN for internal API load balancer target group.
- 22 Specify the **InternalApiTargetGroupArn** value from the output of the CloudFormation template for DNS and load balancing. Use **arn:aws-us-gov** if deploying the cluster to an AWS GovCloud region.
- 23 The ARN for internal service load balancer target group.
- 24 Specify the **InternalServiceTargetGroupArn** value from the output of the CloudFormation template for DNS and load balancing. Use **arn:aws-us-gov** if deploying the cluster to an AWS GovCloud region.

5. Copy the template from the **CloudFormation template for the bootstrap machines** section of this topic and save it as a YAML file on your computer. This template describes the bootstrap machine that your cluster requires.
6. Optional: If you are deploying the cluster with a proxy, you must update the ignition in the template to add the **ignition.config.proxy** fields. Additionally, if you have added the Amazon EC2, Elastic Load Balancing, and S3 VPC endpoints to your VPC, you must add these endpoints to the **noProxy** field.
7. Launch the CloudFormation template to create a stack of AWS resources that represent the bootstrap node:



IMPORTANT

You must enter the command on a single line.

```
$ aws cloudformation create-stack --stack-name <name> ❶
  --template-body file://<template>.yaml ❷
  --parameters file://<parameters>.json ❸
  --capabilities CAPABILITY_NAMED_IAM ❹
```

- ❶ **<name>** is the name for the CloudFormation stack, such as **cluster-bootstrap**. You need the name of this stack if you remove the cluster.
- ❷ **<template>** is the relative path to and name of the CloudFormation template YAML file that you saved.
- ❸ **<parameters>** is the relative path to and name of the CloudFormation parameters JSON file.
- ❹ You must explicitly declare the **CAPABILITY_NAMED_IAM** capability because the provided template creates some **AWS::IAM::Role** and **AWS::IAM::InstanceProfile** resources.

Example output

```
arn:aws:cloudformation:us-east-1:269333783861:stack/cluster-bootstrap/12944486-2add-11eb-9dee-12dace8e3a83
```

8. Confirm that the template components exist:

```
$ aws cloudformation describe-stacks --stack-name <name>
```

After the **StackStatus** displays **CREATE_COMPLETE**, the output displays values for the following parameters. You must provide these parameter values to the other CloudFormation templates that you run to create your cluster:

Bootstrap InstanceId	The bootstrap Instance ID.
-----------------------------	----------------------------

Bootstrap PublicIp	The bootstrap node public IP address.
Bootstrap PrivateIp	The bootstrap node private IP address.

6.4.3.9.1. CloudFormation template for the bootstrap machine

You can use the following CloudFormation template to deploy the bootstrap machine that you need for your OpenShift Container Platform cluster.

Example 6.75. CloudFormation template for the bootstrap machine

```

AWSTemplateFormatVersion: 2010-09-09
Description: Template for OpenShift Cluster Bootstrap (EC2 Instance, Security Groups and IAM)

Parameters:
  InfrastructureName:
    AllowedPattern: ^([a-zA-Z][a-zA-Z0-9\-\_]{0,26})$
    MaxLength: 27
    MinLength: 1
    ConstraintDescription: Infrastructure name must be alphanumeric, start with a letter, and have a
    maximum of 27 characters.
    Description: A short, unique cluster ID used to tag cloud resources and identify items owned or
    used by the cluster.
    Type: String
  RhcosAmi:
    Description: Current Red Hat Enterprise Linux CoreOS AMI to use for bootstrap.
    Type: AWS::EC2::Image::Id
  AllowedBootstrapSshCidr:
    AllowedPattern: ^((([0-9]{1,3}[0-9]{1,3}|[0-9]{1,3}[0-9]{1,3}|[0-9]{1,3}[0-9]{1,3}|[0-9]{1,3}[0-9]{1,3})
    |([0-9]{1,3}[0-9]{1,3}|[0-9]{1,3}[0-9]{1,3}|[0-9]{1,3}[0-9]{1,3}|[0-9]{1,3}[0-9]{1,3})|([0-9]{1,3}[0-9]{1,3}|[0-9]{1,3}[0-9]{1,3}|[0-9]{1,3}[0-9]{1,3}|[0-9]{1,3}[0-9]{1,3})|([0-9]{1,3}[0-9]{1,3}|[0-9]{1,3}[0-9]{1,3}|[0-9]{1,3}[0-9]{1,3}|[0-9]{1,3}[0-9]{1,3}))$
    ConstraintDescription: CIDR block parameter must be in the form x.x.x.x/0-32.
    Default: 0.0.0.0/0
    Description: CIDR block to allow SSH access to the bootstrap node.
    Type: String
  PublicSubnet:
    Description: The public subnet to launch the bootstrap node into.
    Type: AWS::EC2::Subnet::Id
  MasterSecurityGroupId:
    Description: The master security group ID for registering temporary rules.
    Type: AWS::EC2::SecurityGroup::Id
  VpcId:
    Description: The VPC-scoped resources will belong to this VPC.
    Type: AWS::EC2::VPC::Id
  BootstrapIgnitionLocation:
    Default: s3://my-s3-bucket/bootstrap.ign
    Description: Ignition config file location.
    Type: String
  AutoRegisterELB:
    Default: "yes"
    AllowedValues:
      - "yes"
      - "no"

```

Description: Do you want to invoke NLB registration, which requires a Lambda ARN parameter?

Type: String

RegisterNlbTargetsLambdaArn:

Description: ARN for NLB IP target registration lambda.

Type: String

ExternalApiTargetGroupArn:

Description: ARN for external API load balancer target group.

Type: String

InternalApiTargetGroupArn:

Description: ARN for internal API load balancer target group.

Type: String

InternalServiceTargetGroupArn:

Description: ARN for internal service load balancer target group.

Type: String

BootstrapInstanceType:

Description: Instance type for the bootstrap EC2 instance

Default: "i3.large"

Type: String

Metadata:

AWS::CloudFormation::Interface:

ParameterGroups:

- Label:

default: "Cluster Information"

Parameters:

- InfrastructureName

- Label:

default: "Host Information"

Parameters:

- RhcosAmi

- BootstrapIgnitionLocation

- MasterSecurityGroupId

- Label:

default: "Network Configuration"

Parameters:

- VpcId

- AllowedBootstrapSshCidr

- PublicSubnet

- Label:

default: "Load Balancer Automation"

Parameters:

- AutoRegisterELB

- RegisterNlbTargetsLambdaArn

- ExternalApiTargetGroupArn

- InternalApiTargetGroupArn

- InternalServiceTargetGroupArn

ParameterLabels:

InfrastructureName:

default: "Infrastructure Name"

VpcId:

default: "VPC ID"

AllowedBootstrapSshCidr:

default: "Allowed SSH Source"

PublicSubnet:

default: "Public Subnet"

RhcosAmi:

```
    default: "Red Hat Enterprise Linux CoreOS AMI ID"
BootstrapIgnitionLocation:
    default: "Bootstrap Ignition Source"
MasterSecurityGroupId:
    default: "Master Security Group ID"
AutoRegisterELB:
    default: "Use Provided ELB Automation"
```

Conditions:

```
DoRegistration: !Equals ["yes", !Ref AutoRegisterELB]
```

Resources:

```
BootstrapIamRole:
  Type: AWS::IAM::Role
  Properties:
    AssumeRolePolicyDocument:
      Version: "2012-10-17"
      Statement:
        - Effect: "Allow"
          Principal:
            Service:
              - "ec2.amazonaws.com"
          Action:
            - "sts:AssumeRole"
    Path: "/"
  Policies:
    - PolicyName: !Join ["-", [!Ref InfrastructureName, "bootstrap", "policy"]]
      PolicyDocument:
        Version: "2012-10-17"
        Statement:
          - Effect: "Allow"
            Action: "ec2:Describe*"
            Resource: "*"
          - Effect: "Allow"
            Action: "ec2:AttachVolume"
            Resource: "*"
          - Effect: "Allow"
            Action: "ec2:DetachVolume"
            Resource: "*"
          - Effect: "Allow"
            Action: "s3:GetObject"
            Resource: "*"

```

```
BootstrapInstanceProfile:
  Type: "AWS::IAM::InstanceProfile"
  Properties:
    Path: "/"
    Roles:
      - Ref: "BootstrapIamRole"
```

```
BootstrapSecurityGroup:
  Type: AWS::EC2::SecurityGroup
  Properties:
    GroupDescription: Cluster Bootstrap Security Group
    SecurityGroupIngress:
      - IpProtocol: tcp
```

```

FromPort: 22
ToPort: 22
CidrIp: !Ref AllowedBootstrapSshCidr
- IpProtocol: tcp
  ToPort: 19531
  FromPort: 19531
  CidrIp: 0.0.0.0/0
VpCid: !Ref VpCid

```

BootstrapInstance:

```

Type: AWS::EC2::Instance
Properties:
  ImageId: !Ref RhcosAmi
  IamInstanceProfile: !Ref BootstrapInstanceProfile
  InstanceType: !Ref BootstrapInstanceType
  NetworkInterfaces:
  - AssociatePublicIpAddress: "true"
    DeviceIndex: "0"
    GroupSet:
    - !Ref "BootstrapSecurityGroup"
    - !Ref "MasterSecurityGroup"
    SubnetId: !Ref "PublicSubnet"
  UserData:
    Fn::Base64: !Sub
      - '{"ignition":{"config":{"replace":{"source":"${S3Loc}"},"version":"3.1.0"}}}'
      - {
        S3Loc: !Ref BootstrapIgnitionLocation
      }

```

RegisterBootstrapApiTarget:

```

Condition: DoRegistration
Type: Custom::NLBRegister
Properties:
  ServiceToken: !Ref RegisterNLBTargetsLambdaArn
  TargetArn: !Ref ExternalApiTargetGroupArn
  TargetIp: !GetAtt BootstrapInstance.PrivateIp

```

RegisterBootstrapInternalApiTarget:

```

Condition: DoRegistration
Type: Custom::NLBRegister
Properties:
  ServiceToken: !Ref RegisterNLBTargetsLambdaArn
  TargetArn: !Ref InternalApiTargetGroupArn
  TargetIp: !GetAtt BootstrapInstance.PrivateIp

```

RegisterBootstrapInternalServiceTarget:

```

Condition: DoRegistration
Type: Custom::NLBRegister
Properties:
  ServiceToken: !Ref RegisterNLBTargetsLambdaArn
  TargetArn: !Ref InternalServiceTargetGroupArn
  TargetIp: !GetAtt BootstrapInstance.PrivateIp

```

Outputs:

```

BootstrapInstanceCid:
  Description: Bootstrap Instance ID.

```

Value: !Ref BootstrapInstance

BootstrapPublicIp:

Description: The bootstrap node public IP address.

Value: !GetAtt BootstrapInstance.PublicIp

BootstrapPrivateIp:

Description: The bootstrap node private IP address.

Value: !GetAtt BootstrapInstance.PrivateIp


Additional resources

- You can view details about the CloudFormation stacks that you create by navigating to the [AWS CloudFormation console](#).
- See [RHCOS AMIs for the AWS infrastructure](#) for details about the Red Hat Enterprise Linux CoreOS (RHCOS) AMIs for the AWS zones.

6.4.3.10. Creating the control plane machines in AWS


You must create the control plane machines in Amazon Web Services (AWS) that your cluster will use.

You can use the provided CloudFormation template and a custom parameter file to create a stack of AWS resources that represent the control plane nodes.



IMPORTANT

The CloudFormation template creates a stack that represents three control plane nodes.



NOTE

If you do not use the provided CloudFormation template to create your control plane nodes, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- You configured an AWS account.
- You added your AWS keys and region to your local AWS profile by running **aws configure**.
- You generated the Ignition config files for your cluster.
- You created and configured a VPC and associated subnets in AWS.
- You created and configured DNS, load balancers, and listeners in AWS.
- You created the security groups and roles required for your cluster in AWS.
- You created the bootstrap machine.

Procedure

1. Create a JSON file that contains the parameter values that the template requires:

```
[
  {
    "ParameterKey": "InfrastructureName", 1
    "ParameterValue": "mycluster-<random_string>" 2
  },
  {
    "ParameterKey": "RhcocAmi", 3
    "ParameterValue": "ami-<random_string>" 4
  },
  {
    "ParameterKey": "AutoRegisterDNS", 5
    "ParameterValue": "yes" 6
  },
  {
    "ParameterKey": "PrivateHostedZoneId", 7
    "ParameterValue": "<random_string>" 8
  },
  {
    "ParameterKey": "PrivateHostedZoneName", 9
    "ParameterValue": "mycluster.example.com" 10
  },
  {
    "ParameterKey": "Master0Subnet", 11
    "ParameterValue": "subnet-<random_string>" 12
  },
  {
    "ParameterKey": "Master1Subnet", 13
    "ParameterValue": "subnet-<random_string>" 14
  },
  {
    "ParameterKey": "Master2Subnet", 15
    "ParameterValue": "subnet-<random_string>" 16
  },
  {
    "ParameterKey": "MasterSecurityGroupID", 17
    "ParameterValue": "sg-<random_string>" 18
  },
  {
    "ParameterKey": "IgnitionLocation", 19
    "ParameterValue": "https://api-int.<cluster_name>.<domain_name>:22623/config/master"
    20
  },
  {
    "ParameterKey": "CertificateAuthorities", 21
    "ParameterValue": "data:text/plain;charset=utf-8;base64,ABC...xYz==" 22
  },
  {
    "ParameterKey": "MasterInstanceProfileName", 23
    "ParameterValue": "<roles_stack>-MasterInstanceProfile-<random_string>" 24
  },
  {

```

```

    "ParameterKey": "MasterInstanceType", 25
    "ParameterValue": "" 26
  },
  {
    "ParameterKey": "AutoRegisterELB", 27
    "ParameterValue": "yes" 28
  },
  {
    "ParameterKey": "RegisterNlbTargetsLambdaArn", 29
    "ParameterValue": "arn:aws:lambda:<aws_region>:<account_number>:function:
<dns_stack_name>-RegisterNlbTargets-<random_string>" 30
  },
  {
    "ParameterKey": "ExternalApiTargetGroupArn", 31
    "ParameterValue": "arn:aws:elasticloadbalancing:<aws_region>:
<account_number>:targetgroup/<dns_stack_name>-Exter-<random_string>" 32
  },
  {
    "ParameterKey": "InternalApiTargetGroupArn", 33
    "ParameterValue": "arn:aws:elasticloadbalancing:<aws_region>:
<account_number>:targetgroup/<dns_stack_name>-Inter-<random_string>" 34
  },
  {
    "ParameterKey": "InternalServiceTargetGroupArn", 35
    "ParameterValue": "arn:aws:elasticloadbalancing:<aws_region>:
<account_number>:targetgroup/<dns_stack_name>-Inter-<random_string>" 36
  }
]

```

- 1 The name for your cluster infrastructure that is encoded in your Ignition config files for the cluster.
- 2 Specify the infrastructure name that you extracted from the Ignition config file metadata, which has the format **<cluster-name>-<random-string>**.
- 3 Current Red Hat Enterprise Linux CoreOS (RHCOS) AMI to use for the control plane machines based on your selected architecture.
- 4 Specify an **AWS::EC2::Image::Id** value.
- 5 Whether or not to perform DNS etcd registration.
- 6 Specify **yes** or **no**. If you specify **yes**, you must provide hosted zone information.
- 7 The Route 53 private zone ID to register the etcd targets with.
- 8 Specify the **PrivateHostedZoneId** value from the output of the CloudFormation template for DNS and load balancing.
- 9 The Route 53 zone to register the targets with.
- 10 Specify **<cluster_name>.<domain_name>** where **<domain_name>** is the Route 53 base domain that you used when you generated **install-config.yaml** file for the cluster. Do not include the trailing period (.) that is displayed in the AWS console.

- 11 13 15 A subnet, preferably private, to launch the control plane machines on.
- 12 14 16 Specify a subnet from the **PrivateSubnets** value from the output of the CloudFormation template for DNS and load balancing.
- 17 The master security group ID to associate with control plane nodes.
- 18 Specify the **MasterSecurityGroupId** value from the output of the CloudFormation template for the security group and roles.
- 19 The location to fetch control plane Ignition config file from.
- 20 Specify the generated Ignition config file location, https://api-int.<cluster_name>.<domain_name>:22623/config/master.
- 21 The base64 encoded certificate authority string to use.
- 22 Specify the value from the **master.ign** file that is in the installation directory. This value is the long string with the format **data:text/plain;charset=utf-8;base64,ABC...xYz==**.
- 23 The IAM profile to associate with control plane nodes.
- 24 Specify the **MasterInstanceProfile** parameter value from the output of the CloudFormation template for the security group and roles.
- 25 The type of AWS instance to use for the control plane machines based on your selected architecture.
- 26 The instance type value corresponds to the minimum resource requirements for control plane machines. For example **m6i.xlarge** is a type for AMD64 and **m6g.xlarge** is a type for ARM64.
- 27 Whether or not to register a network load balancer (NLB).
- 28 Specify **yes** or **no**. If you specify **yes**, you must provide a Lambda Amazon Resource Name (ARN) value.
- 29 The ARN for NLB IP target registration lambda group.
- 30 Specify the **RegisterNlbTargetsLambda** value from the output of the CloudFormation template for DNS and load balancing. Use **arn:aws-us-gov** if deploying the cluster to an AWS GovCloud region.
- 31 The ARN for external API load balancer target group.
- 32 Specify the **ExternalApiTargetGroupArn** value from the output of the CloudFormation template for DNS and load balancing. Use **arn:aws-us-gov** if deploying the cluster to an AWS GovCloud region.
- 33 The ARN for internal API load balancer target group.
- 34 Specify the **InternalApiTargetGroupArn** value from the output of the CloudFormation template for DNS and load balancing. Use **arn:aws-us-gov** if deploying the cluster to an AWS GovCloud region.
- 35 The ARN for internal service load balancer target group.

- 36 Specify the **InternalServiceTargetGroupArn** value from the output of the CloudFormation template for DNS and load balancing. Use **arn:aws-us-gov** if deploying
2. Copy the template from the **CloudFormation template for control plane machines** section of this topic and save it as a YAML file on your computer. This template describes the control plane machines that your cluster requires.
3. If you specified an **m5** instance type as the value for **MasterInstanceType**, add that instance type to the **MasterInstanceType.AllowedValues** parameter in the CloudFormation template.
4. Launch the CloudFormation template to create a stack of AWS resources that represent the control plane nodes:

**IMPORTANT**

You must enter the command on a single line.

```
$ aws cloudformation create-stack --stack-name <name> 1
  --template-body file://<template>.yaml 2
  --parameters file://<parameters>.json 3
```

- 1 **<name>** is the name for the CloudFormation stack, such as **cluster-control-plane**. You need the name of this stack if you remove the cluster.
- 2 **<template>** is the relative path to and name of the CloudFormation template YAML file that you saved.
- 3 **<parameters>** is the relative path to and name of the CloudFormation parameters JSON file.

Example output

```
arn:aws:cloudformation:us-east-1:269333783861:stack/cluster-control-plane/21c7e2b0-2ee2-11eb-c6f6-0aa34627df4b
```

**NOTE**

The CloudFormation template creates a stack that represents three control plane nodes.

5. Confirm that the template components exist:

```
$ aws cloudformation describe-stacks --stack-name <name>
```

6.4.3.10.1. CloudFormation template for control plane machines

You can use the following CloudFormation template to deploy the control plane machines that you need for your OpenShift Container Platform cluster.

Example 6.76. CloudFormation template for control plane machines

AWSTemplateFormatVersion: 2010-09-09

Description: Template for OpenShift Cluster Node Launch (EC2 master instances)

Parameters:

InfrastructureName:

AllowedPattern: ^([a-zA-Z][a-zA-Z0-9\-_]{0,26})\$

MaxLength: 27

MinLength: 1

ConstraintDescription: Infrastructure name must be alphanumeric, start with a letter, and have a maximum of 27 characters.

Description: A short, unique cluster ID used to tag nodes for the kubelet cloud provider.

Type: String

RhcosAmi:

Description: Current Red Hat Enterprise Linux CoreOS AMI to use for bootstrap.

Type: AWS::EC2::Image::Id

AutoRegisterDNS:

Default: ""

Description: unused

Type: String

PrivateHostedZoneId:

Default: ""

Description: unused

Type: String

PrivateHostedZoneName:

Default: ""

Description: unused

Type: String

Master0Subnet:

Description: The subnets, recommend private, to launch the master nodes into.

Type: AWS::EC2::Subnet::Id

Master1Subnet:

Description: The subnets, recommend private, to launch the master nodes into.

Type: AWS::EC2::Subnet::Id

Master2Subnet:

Description: The subnets, recommend private, to launch the master nodes into.

Type: AWS::EC2::Subnet::Id

MasterSecurityGroupId:

Description: The master security group ID to associate with master nodes.

Type: AWS::EC2::SecurityGroup::Id

IgnitionLocation:

Default: https://api-int.\$CLUSTER_NAME.\$DOMAIN:22623/config/master

Description: Ignition config file location.

Type: String

CertificateAuthorities:

Default: data:text/plain;charset=utf-8;base64,ABC...xYz==

Description: Base64 encoded certificate authority string to use.

Type: String

MasterInstanceProfileName:

Description: IAM profile to associate with master nodes.

Type: String

MasterInstanceType:

Default: m5.xlarge

Type: String

AutoRegisterELB:

Default: "yes"

AllowedValues:

- "yes"
- "no"

Description: Do you want to invoke NLB registration, which requires a Lambda ARN parameter?

Type: String

RegisterNlbTargetsLambdaArn:

Description: ARN for NLB IP target registration lambda. Supply the value from the cluster infrastructure or select "no" for AutoRegisterELB.

Type: String

ExternalApiTargetGroupArn:

Description: ARN for external API load balancer target group. Supply the value from the cluster infrastructure or select "no" for AutoRegisterELB.

Type: String

InternalApiTargetGroupArn:

Description: ARN for internal API load balancer target group. Supply the value from the cluster infrastructure or select "no" for AutoRegisterELB.

Type: String

InternalServiceTargetGroupArn:

Description: ARN for internal service load balancer target group. Supply the value from the cluster infrastructure or select "no" for AutoRegisterELB.

Type: String

Metadata:

AWS::CloudFormation::Interface:

ParameterGroups:

- Label:

default: "Cluster Information"

Parameters:

- InfrastructureName

- Label:

default: "Host Information"

Parameters:

- MasterInstanceType

- RhcosAmi

- IgnitionLocation

- CertificateAuthorities

- MasterSecurityGroupId

- MasterInstanceProfileName

- Label:

default: "Network Configuration"

Parameters:

- VpcId

- AllowedBootstrapSshCidr

- Master0Subnet

- Master1Subnet

- Master2Subnet

- Label:

default: "Load Balancer Automation"

Parameters:

- AutoRegisterELB

- RegisterNlbTargetsLambdaArn

- ExternalApiTargetGroupArn

- InternalApiTargetGroupArn

- InternalServiceTargetGroupArn

ParameterLabels:

InfrastructureName:

```

    default: "Infrastructure Name"
VpcId:
    default: "VPC ID"
Master0Subnet:
    default: "Master-0 Subnet"
Master1Subnet:
    default: "Master-1 Subnet"
Master2Subnet:
    default: "Master-2 Subnet"
MasterInstanceType:
    default: "Master Instance Type"
MasterInstanceProfileName:
    default: "Master Instance Profile Name"
RhcOsAmi:
    default: "Red Hat Enterprise Linux CoreOS AMI ID"
BootstrapIgnitionLocation:
    default: "Master Ignition Source"
CertificateAuthorities:
    default: "Ignition CA String"
MasterSecurityGroupId:
    default: "Master Security Group ID"
AutoRegisterELB:
    default: "Use Provided ELB Automation"

```

Conditions:

```
DoRegistration: !Equals ["yes", !Ref AutoRegisterELB]
```

Resources:

Master0:

```
Type: AWS::EC2::Instance
```

Properties:

```
ImageId: !Ref RhcOsAmi
```

```
BlockDeviceMappings:
```

```
- DeviceName: /dev/xvda
```

```
Ebs:
```

```
VolumeSize: "120"
```

```
VolumeType: "gp2"
```

```
IamInstanceProfile: !Ref MasterInstanceProfileName
```

```
InstanceType: !Ref MasterInstanceType
```

```
NetworkInterfaces:
```

```
- AssociatePublicIpAddress: "false"
```

```
DeviceIndex: "0"
```

```
GroupSet:
```

```
- !Ref "MasterSecurityGroupId"
```

```
SubnetId: !Ref "Master0Subnet"
```

```
UserData:
```

```
Fn::Base64: !Sub
```

```
- {"ignition":{"config":{"merge":{"source":"${SOURCE}"},"security":{"tls":
{"certificateAuthorities":[{"source":"${CA_BUNDLE}"]}},"version":"3.1.0"}}
```

```
- {
```

```
  SOURCE: !Ref IgnitionLocation,
```

```
  CA_BUNDLE: !Ref CertificateAuthorities,
```

```
}
```

```
Tags:
```

```
- Key: !Join [ "", ["kubernetes.io/cluster/", !Ref InfrastructureName]]
```

```
Value: "shared"
```

RegisterMaster0:

Condition: DoRegistration

Type: Custom::NLBRegister

Properties:

ServiceToken: !Ref RegisterNlbTargetsLambdaArn

TargetArn: !Ref ExternalApiTargetGroupArn

TargetIp: !GetAtt Master0.PrivateIp

RegisterMaster0InternalApiTarget:

Condition: DoRegistration

Type: Custom::NLBRegister

Properties:

ServiceToken: !Ref RegisterNlbTargetsLambdaArn

TargetArn: !Ref InternalApiTargetGroupArn

TargetIp: !GetAtt Master0.PrivateIp

RegisterMaster0InternalServiceTarget:

Condition: DoRegistration

Type: Custom::NLBRegister

Properties:

ServiceToken: !Ref RegisterNlbTargetsLambdaArn

TargetArn: !Ref InternalServiceTargetGroupArn

TargetIp: !GetAtt Master0.PrivateIp

Master1:

Type: AWS::EC2::Instance

Properties:

ImageId: !Ref RhcosAmi

BlockDeviceMappings:

- DeviceName: /dev/xvda

Ebs:

VolumeSize: "120"

VolumeType: "gp2"

IamInstanceProfile: !Ref MasterInstanceProfileName

InstanceType: !Ref MasterInstanceType

NetworkInterfaces:

- AssociatePublicIpAddress: "false"

DeviceIndex: "0"

GroupSet:

- !Ref "MasterSecurityGroupId"

SubnetId: !Ref "Master1Subnet"

UserData:

Fn::Base64: !Sub

```
- {"ignition":{"config":{"merge":{"source":"${SOURCE}"},"security":{"tls":{"certificateAuthorities":[{"source":"${CA_BUNDLE}"]}},"version":"3.1.0"}}
```

- {

SOURCE: !Ref IgnitionLocation,

CA_BUNDLE: !Ref CertificateAuthorities,

}

Tags:

- Key: !Join ["", ["kubernetes.io/cluster/", !Ref InfrastructureName]]

Value: "shared"

RegisterMaster1:

Condition: DoRegistration

Type: Custom::NLBRegister
 Properties:
 ServiceToken: !Ref RegisterNlbTargetsLambdaArn
 TargetArn: !Ref ExternalApiTargetGroupArn
 TargetIp: !GetAtt Master1.PrivateIp

RegisterMaster1InternalApiTarget:
 Condition: DoRegistration
 Type: Custom::NLBRegister
 Properties:
 ServiceToken: !Ref RegisterNlbTargetsLambdaArn
 TargetArn: !Ref InternalApiTargetGroupArn
 TargetIp: !GetAtt Master1.PrivateIp

RegisterMaster1InternalServiceTarget:
 Condition: DoRegistration
 Type: Custom::NLBRegister
 Properties:
 ServiceToken: !Ref RegisterNlbTargetsLambdaArn
 TargetArn: !Ref InternalServiceTargetGroupArn
 TargetIp: !GetAtt Master1.PrivateIp

Master2:
 Type: AWS::EC2::Instance
 Properties:
 ImageId: !Ref RhcosAmi
 BlockDeviceMappings:
 - DeviceName: /dev/xvda
 Ebs:
 VolumeSize: "120"
 VolumeType: "gp2"
 IamInstanceProfile: !Ref MasterInstanceProfileName
 InstanceType: !Ref MasterInstanceType
 NetworkInterfaces:
 - AssociatePublicIpAddress: "false"
 DeviceIndex: "0"
 GroupSet:
 - !Ref "MasterSecurityGroupId"
 SubnetId: !Ref "Master2Subnet"
 UserData:
 Fn::Base64: !Sub
 - '{"ignition":{"config":{"merge":[{"source":"\${SOURCE}"}]},"security":{"tls":
 {"certificateAuthorities":[{"source":"\${CA_BUNDLE}"}]},"version":"3.1.0"}}'
 - {
 SOURCE: !Ref IgnitionLocation,
 CA_BUNDLE: !Ref CertificateAuthorities,
 }
 Tags:
 - Key: !Join ["", ["kubernetes.io/cluster/", !Ref InfrastructureName]]
 Value: "shared"

RegisterMaster2:
 Condition: DoRegistration
 Type: Custom::NLBRegister
 Properties:
 ServiceToken: !Ref RegisterNlbTargetsLambdaArn

```
TargetArn: !Ref ExternalApiTargetGroupArn
TargetIp: !GetAtt Master2.PrivateIp
```

```
RegisterMaster2InternalApiTarget:
Condition: DoRegistration
Type: Custom::NLBRegister
Properties:
ServiceToken: !Ref RegisterNlbTargetsLambdaArn
TargetArn: !Ref InternalApiTargetGroupArn
TargetIp: !GetAtt Master2.PrivateIp
```

```
RegisterMaster2InternalServiceTarget:
Condition: DoRegistration
Type: Custom::NLBRegister
Properties:
ServiceToken: !Ref RegisterNlbTargetsLambdaArn
TargetArn: !Ref InternalServiceTargetGroupArn
TargetIp: !GetAtt Master2.PrivateIp
```

Outputs:

```
PrivateIPs:
Description: The control-plane node private IP addresses.
Value:
!Join [
  ",",
  [!GetAtt Master0.PrivateIp, !GetAtt Master1.PrivateIp, !GetAtt Master2.PrivateIp]
]
```

Additional resources

- You can view details about the CloudFormation stacks that you create by navigating to the [AWS CloudFormation console](#).

6.4.3.11. Creating the worker nodes in AWS

You can create worker nodes in Amazon Web Services (AWS) for your cluster to use.



NOTE

If you are installing a three-node cluster, skip this step. A three-node cluster consists of three control plane machines, which also act as compute machines.

You can use the provided CloudFormation template and a custom parameter file to create a stack of AWS resources that represent a worker node.



IMPORTANT

The CloudFormation template creates a stack that represents one worker node. You must create a stack for each worker node.



NOTE

If you do not use the provided CloudFormation template to create your worker nodes, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- You configured an AWS account.
- You added your AWS keys and region to your local AWS profile by running **aws configure**.
- You generated the Ignition config files for your cluster.
- You created and configured a VPC and associated subnets in AWS.
- You created and configured DNS, load balancers, and listeners in AWS.
- You created the security groups and roles required for your cluster in AWS.
- You created the bootstrap machine.
- You created the control plane machines.

Procedure

1. Create a JSON file that contains the parameter values that the CloudFormation template requires:

```
[
  {
    "ParameterKey": "InfrastructureName", 1
    "ParameterValue": "mycluster-<random_string>" 2
  },
  {
    "ParameterKey": "RhcocAmi", 3
    "ParameterValue": "ami-<random_string>" 4
  },
  {
    "ParameterKey": "Subnet", 5
    "ParameterValue": "subnet-<random_string>" 6
  },
  {
    "ParameterKey": "WorkerSecurityGroupld", 7
    "ParameterValue": "sg-<random_string>" 8
  },
  {
    "ParameterKey": "IgnitionLocation", 9
    "ParameterValue": "https://api-int.<cluster_name>.<domain_name>:22623/config/worker"
  },
  {
    "ParameterKey": "CertificateAuthorities", 11
    "ParameterValue": "" 12
  }
]
```

```

    },
    {
      "ParameterKey": "WorkerInstanceProfileName", 13
      "ParameterValue": "" 14
    },
    {
      "ParameterKey": "WorkerInstanceType", 15
      "ParameterValue": "" 16
    }
  ]

```

- 1 The name for your cluster infrastructure that is encoded in your Ignition config files for the cluster.
 - 2 Specify the infrastructure name that you extracted from the Ignition config file metadata, which has the format **<cluster-name>-<random-string>**.
 - 3 Current Red Hat Enterprise Linux CoreOS (RHCOS) AMI to use for the worker nodes based on your selected architecture.
 - 4 Specify an **AWS::EC2::Image::Id** value.
 - 5 A subnet, preferably private, to start the worker nodes on.
 - 6 Specify a subnet from the **PrivateSubnets** value from the output of the CloudFormation template for DNS and load balancing.
 - 7 The worker security group ID to associate with worker nodes.
 - 8 Specify the **WorkerSecurityGroupId** value from the output of the CloudFormation template for the security group and roles.
 - 9 The location to fetch the bootstrap Ignition config file from.
 - 10 Specify the generated Ignition config location, https://api-int.<cluster_name>.<domain_name>:22623/config/worker.
 - 11 Base64 encoded certificate authority string to use.
 - 12 Specify the value from the **worker.ign** file that is in the installation directory. This value is the long string with the format **data:text/plain;charset=utf-8;base64,ABC...xYz==**.
 - 13 The IAM profile to associate with worker nodes.
 - 14 Specify the **WorkerInstanceProfile** parameter value from the output of the CloudFormation template for the security group and roles.
 - 15 The type of AWS instance to use for the compute machines based on your selected architecture.
 - 16 The instance type value corresponds to the minimum resource requirements for compute machines. For example **m6i.large** is a type for AMD64 and **m6g.large** is a type for ARM64.
2. Copy the template from the **CloudFormation template for worker machines** section of this topic and save it as a YAML file on your computer. This template describes the networking objects and load balancers that your cluster requires.

3. Optional: If you specified an **m5** instance type as the value for **WorkerInstanceType**, add that instance type to the **WorkerInstanceType.AllowedValues** parameter in the CloudFormation template.
4. Optional: If you are deploying with an AWS Marketplace image, update the **Worker0.type.properties.ImageID** parameter with the AMI ID that you obtained from your subscription.
5. Use the CloudFormation template to create a stack of AWS resources that represent a worker node:



IMPORTANT

You must enter the command on a single line.

```
$ aws cloudformation create-stack --stack-name <name> ❶
  --template-body file://<template>.yaml \ ❷
  --parameters file://<parameters>.json ❸
```

- ❶ **<name>** is the name for the CloudFormation stack, such as **cluster-worker-1**. You need the name of this stack if you remove the cluster.
- ❷ **<template>** is the relative path to and name of the CloudFormation template YAML file that you saved.
- ❸ **<parameters>** is the relative path to and name of the CloudFormation parameters JSON file.

Example output

```
arn:aws:cloudformation:us-east-1:269333783861:stack/cluster-worker-1/729ee301-1c2a-
11eb-348f-sd9888c65b59
```



NOTE

The CloudFormation template creates a stack that represents one worker node.

6. Confirm that the template components exist:

```
$ aws cloudformation describe-stacks --stack-name <name>
```

7. Continue to create worker stacks until you have created enough worker machines for your cluster. You can create additional worker stacks by referencing the same template and parameter files and specifying a different stack name.



IMPORTANT

You must create at least two worker machines, so you must create at least two stacks that use this CloudFormation template.

6.4.3.11.1. CloudFormation template for worker machines

You can use the following CloudFormation template to deploy the worker machines that you need for your OpenShift Container Platform cluster.

Example 6.77. CloudFormation template for worker machines

AWSTemplateFormatVersion: 2010-09-09

Description: Template for OpenShift Cluster Node Launch (EC2 worker instance)

Parameters:

InfrastructureName:

AllowedPattern: ^([a-zA-Z][a-zA-Z0-9\-\]{0,26})\$

MaxLength: 27

MinLength: 1

ConstraintDescription: Infrastructure name must be alphanumeric, start with a letter, and have a maximum of 27 characters.

Description: A short, unique cluster ID used to tag nodes for the kubelet cloud provider.

Type: String

RhcosAmi:

Description: Current Red Hat Enterprise Linux CoreOS AMI to use for bootstrap.

Type: AWS::EC2::Image::Id

Subnet:

Description: The subnets, recommend private, to launch the worker nodes into.

Type: AWS::EC2::Subnet::Id

WorkerSecurityGroupId:

Description: The worker security group ID to associate with worker nodes.

Type: AWS::EC2::SecurityGroup::Id

IgnitionLocation:

Default: https://api-int.\$CLUSTER_NAME.\$DOMAIN:22623/config/worker

Description: Ignition config file location.

Type: String

CertificateAuthorities:

Default: data:text/plain;charset=utf-8;base64,ABC...xYz==

Description: Base64 encoded certificate authority string to use.

Type: String

WorkerInstanceProfileName:

Description: IAM profile to associate with worker nodes.

Type: String

WorkerInstanceType:

Default: m5.large

Type: String

Metadata:

AWS::CloudFormation::Interface:

ParameterGroups:

- Label:

default: "Cluster Information"

Parameters:

- InfrastructureName

- Label:

default: "Host Information"

Parameters:

- WorkerInstanceType

- RhcosAmi

- IgnitionLocation

- CertificateAuthorities

- WorkerSecurityGroupId

```

- WorkerInstanceProfileName
- Label:
  default: "Network Configuration"
Parameters:
- Subnet
ParameterLabels:
Subnet:
  default: "Subnet"
InfrastructureName:
  default: "Infrastructure Name"
WorkerInstanceType:
  default: "Worker Instance Type"
WorkerInstanceProfileName:
  default: "Worker Instance Profile Name"
RhcOsAmi:
  default: "Red Hat Enterprise Linux CoreOS AMI ID"
IgnitionLocation:
  default: "Worker Ignition Source"
CertificateAuthorities:
  default: "Ignition CA String"
WorkerSecurityGroupId:
  default: "Worker Security Group ID"

Resources:
Worker0:
  Type: AWS::EC2::Instance
  Properties:
    ImageId: !Ref RhcOsAmi
    BlockDeviceMappings:
      - DeviceName: /dev/xvda
        Ebs:
          VolumeSize: "120"
          VolumeType: "gp2"
    IamInstanceProfile: !Ref WorkerInstanceProfileName
    InstanceType: !Ref WorkerInstanceType
    NetworkInterfaces:
      - AssociatePublicIpAddress: "false"
        DeviceIndex: "0"
        GroupSet:
          - !Ref "WorkerSecurityGroupId"
        SubnetId: !Ref "Subnet"
    UserData:
      Fn::Base64: !Sub
        - '{"ignition":{"config":{"merge":[{"source":"${SOURCE}"}]},"security":{"tls":
{"certificateAuthorities":[{"source":"${CA_BUNDLE}"}]},"version":"3.1.0"}}'
        - {
          SOURCE: !Ref IgnitionLocation,
          CA_BUNDLE: !Ref CertificateAuthorities,
        }
    Tags:
      - Key: !Join [ "", ["kubernetes.io/cluster/", !Ref InfrastructureName] ]
        Value: "shared"

Outputs:

```

PrivateIP:

Description: The compute node private IP address.

Value: !GetAtt Worker0.PrivateIp

Additional resources

- You can view details about the CloudFormation stacks that you create by navigating to the [AWS CloudFormation console](#).

6.4.3.12. Initializing the bootstrap sequence on AWS with user-provisioned infrastructure

After you create all of the required infrastructure in Amazon Web Services (AWS), you can start the bootstrap sequence that initializes the OpenShift Container Platform control plane.

Prerequisites

- You configured an AWS account.
- You added your AWS keys and region to your local AWS profile by running **aws configure**.
- You generated the Ignition config files for your cluster.
- You created and configured a VPC and associated subnets in AWS.
- You created and configured DNS, load balancers, and listeners in AWS.
- You created the security groups and roles required for your cluster in AWS.
- You created the bootstrap machine.
- You created the control plane machines.
- You created the worker nodes.

Procedure

1. Change to the directory that contains the installation program and start the bootstrap process that initializes the OpenShift Container Platform control plane:

```
$ ./openshift-install wait-for bootstrap-complete --dir <installation_directory> \ 1
--log-level=info 2
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

- 2** To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Example output

```
INFO Waiting up to 20m0s for the Kubernetes API at
https://api.mycluster.example.com:6443...
INFO API v1.29.4 up
```

```
INFO Waiting up to 30m0s for bootstrapping to complete...
INFO It is now safe to remove the bootstrap resources
INFO Time elapsed: 1s
```

If the command exits without a **FATAL** warning, your OpenShift Container Platform control plane has initialized.



NOTE

After the control plane initializes, it sets up the compute nodes and installs additional services in the form of Operators.

Additional resources

- See [Monitoring installation progress](#) for details about monitoring the installation, bootstrap, and control plane logs as an OpenShift Container Platform installation progresses.
- See [Gathering bootstrap node diagnostic data](#) for information about troubleshooting issues related to the bootstrap process.
- You can view details about the running instances that are created by using the [AWS EC2 console](#).

6.4.3.13. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

6.4.3.14. Approving the certificate signing requests for your machines

When you add machines to a cluster, two pending certificate signing requests (CSRs) are generated for each machine that you added. You must confirm that these CSRs are approved or, if necessary, approve them yourself. The client requests must be approved first, followed by the server requests.

Prerequisites

- You added machines to your cluster.

Procedure

- Confirm that the cluster recognizes the machines:

```
$ oc get nodes
```

Example output

```
NAME      STATUS  ROLES  AGE  VERSION
master-0  Ready   master 63m  v1.29.4
master-1  Ready   master 63m  v1.29.4
master-2  Ready   master 64m  v1.29.4
```

The output lists all of the machines that you created.



NOTE

The preceding output might not include the compute nodes, also known as worker nodes, until some CSRs are approved.

- Review the pending CSRs and ensure that you see the client requests with the **Pending** or **Approved** status for each machine that you added to the cluster:


```
$ oc get csr
```

Example output

```
NAME      AGE  REQUESTOR                                     CONDITION
csr-8b2br 15m  system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper  Pending
csr-8vnps 15m  system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper  Pending
...
```

In this example, two machines are joining the cluster. You might see more approved CSRs in the list.

- If the CSRs were not approved, after all of the pending CSRs for the machines you added are in **Pending** status, approve the CSRs for your cluster machines:

**NOTE**

Because the CSRs rotate automatically, approve your CSRs within an hour of adding the machines to the cluster. If you do not approve them within an hour, the certificates will rotate, and more than two certificates will be present for each node. You must approve all of these certificates. After the client CSR is approved, the Kubelet creates a secondary CSR for the serving certificate, which requires manual approval. Then, subsequent serving certificate renewal requests are automatically approved by the **machine-approver** if the Kubelet requests a new certificate with identical parameters.

**NOTE**

For clusters running on platforms that are not machine API enabled, such as bare metal and other user-provisioned infrastructure, you must implement a method of automatically approving the kubelet serving certificate requests (CSRs). If a request is not approved, then the **oc exec**, **oc rsh**, and **oc logs** commands cannot succeed, because a serving certificate is required when the API server connects to the kubelet. Any operation that contacts the Kubelet endpoint requires this certificate approval to be in place. The method must watch for new CSRs, confirm that the CSR was submitted by the **node-bootstrap** service account in the **system:node** or **system:admin** groups, and confirm the identity of the node.

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{"\n"}{{end}}{{end}}' | xargs --no-run-if-empty oc adm certificate approve
```

**NOTE**

Some Operators might not become available until some CSRs are approved.

4. Now that your client requests are approved, you must review the server requests for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

```
NAME      AGE   REQUESTOR                                     CONDITION
csr-bfd72 5m26s system:node:ip-10-0-50-126.us-east-2.compute.internal
Pending
csr-c57lv 5m26s system:node:ip-10-0-95-157.us-east-2.compute.internal
Pending
...
```

5. If the remaining CSRs are not approved, and are in the **Pending** status, approve the CSRs for your cluster machines:

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

1 **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:


```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{"\n"}\n{{end}}{{end}}' | xargs oc adm certificate approve
```

6. After all client and server CSRs have been approved, the machines have the **Ready** status. Verify this by running the following command:

```
$ oc get nodes
```

Example output

```
NAME      STATUS  ROLES  AGE  VERSION
master-0  Ready   master 73m  v1.29.4
master-1  Ready   master 73m  v1.29.4
master-2  Ready   master 74m  v1.29.4
worker-0  Ready   worker 11m  v1.29.4
worker-1  Ready   worker 11m  v1.29.4
```



NOTE

It can take a few minutes after approval of the server CSRs for the machines to transition to the **Ready** status.

Additional information

- For more information on CSRs, see [Certificate Signing Requests](#).

6.4.3.15. Initial Operator configuration

After the control plane initializes, you must immediately configure some Operators so that they all become available.

Prerequisites

- Your control plane has initialized.

Procedure

- Watch the cluster components come online:

```
$ watch -n5 oc get clusteroperators
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.16.0	True	False	False	19m
baremetal	4.16.0	True	False	False	37m
cloud-credential	4.16.0	True	False	False	40m
cluster-autoscaler	4.16.0	True	False	False	37m
config-operator	4.16.0	True	False	False	38m
console	4.16.0	True	False	False	26m
csi-snapshot-controller	4.16.0	True	False	False	37m
dns	4.16.0	True	False	False	37m
etcd	4.16.0	True	False	False	36m
image-registry	4.16.0	True	False	False	31m
ingress	4.16.0	True	False	False	30m
insights	4.16.0	True	False	False	31m
kube-apiserver	4.16.0	True	False	False	26m
kube-controller-manager	4.16.0	True	False	False	36m
kube-scheduler	4.16.0	True	False	False	36m
kube-storage-version-migrator	4.16.0	True	False	False	37m
machine-api	4.16.0	True	False	False	29m
machine-approver	4.16.0	True	False	False	37m
machine-config	4.16.0	True	False	False	36m
marketplace	4.16.0	True	False	False	37m
monitoring	4.16.0	True	False	False	29m
network	4.16.0	True	False	False	38m
node-tuning	4.16.0	True	False	False	37m
openshift-apiserver	4.16.0	True	False	False	32m
openshift-controller-manager	4.16.0	True	False	False	30m
openshift-samples	4.16.0	True	False	False	32m
operator-lifecycle-manager	4.16.0	True	False	False	37m
operator-lifecycle-manager-catalog	4.16.0	True	False	False	37m
operator-lifecycle-manager-packageserver	4.16.0	True	False	False	32m
service-ca	4.16.0	True	False	False	38m
storage	4.16.0	True	False	False	37m

2. Configure the Operators that are not available.

6.4.3.15.1. Image registry storage configuration

Amazon Web Services provides default storage, which means the Image Registry Operator is available after installation. However, if the Registry Operator cannot create an S3 bucket and automatically configure storage, you must manually configure registry storage.

Instructions are shown for configuring a persistent volume, which is required for production clusters. Where applicable, instructions are shown for configuring an empty directory as the storage location, which is available for only non-production clusters.

Additional instructions are provided for allowing the image registry to use block storage types by using the **Recreate** rollout strategy during upgrades.

You can configure registry storage for user-provisioned infrastructure in AWS to deploy OpenShift Container Platform to hidden regions. See [Configuring the registry for AWS user-provisioned infrastructure](#) for more information.

6.4.3.15.1.1. Configuring registry storage for AWS with user-provisioned infrastructure

During installation, your cloud credentials are sufficient to create an Amazon S3 bucket and the Registry Operator will automatically configure storage.

If the Registry Operator cannot create an S3 bucket and automatically configure storage, you can create an S3 bucket and configure storage with the following procedure.

Prerequisites

- You have a cluster on AWS with user-provisioned infrastructure.
- For Amazon S3 storage, the secret is expected to contain two keys:
 - **REGISTRY_STORAGE_S3_ACCESSKEY**
 - **REGISTRY_STORAGE_S3_SECRETKEY**

Procedure

Use the following procedure if the Registry Operator cannot create an S3 bucket and automatically configure storage.

1. Set up a [Bucket Lifecycle Policy](#) to abort incomplete multipart uploads that are one day old.
2. Fill in the storage configuration in **configs.imageregistry.operator.openshift.io/cluster**:

```
$ oc edit configs.imageregistry.operator.openshift.io/cluster
```

Example configuration

```
storage:
  s3:
    bucket: <bucket-name>
    region: <region-name>
```



WARNING

To secure your registry images in AWS, [block public access](#) to the S3 bucket.

6.4.3.15.1.2. Configuring storage for the image registry in non-production clusters

You must configure storage for the Image Registry Operator. For non-production clusters, you can set the image registry to an empty directory. If you do so, all images are lost if you restart the registry.

Procedure

- To set the image registry storage to an empty directory:

```
$ oc patch configs.imageregistry.operator.openshift.io cluster --type merge --patch '{"spec": {"storage":{"emptyDir":{}}}'
```

■

**WARNING**

Configure this option for only non-production clusters.

If you run this command before the Image Registry Operator initializes its components, the **oc patch** command fails with the following error:

```
Error from server (NotFound): configs.imageregistry.operator.openshift.io "cluster" not found
```

Wait a few minutes and run the command again.

6.4.3.16. Deleting the bootstrap resources

After you complete the initial Operator configuration for the cluster, remove the bootstrap resources from Amazon Web Services (AWS).

Prerequisites

- You completed the initial Operator configuration for your cluster.

Procedure

1. Delete the bootstrap resources. If you used the CloudFormation template, [delete its stack](#):

- Delete the stack by using the AWS CLI:

```
$ aws cloudformation delete-stack --stack-name <name> 1
```

1 **<name>** is the name of your bootstrap stack.

- Delete the stack by using the [AWS CloudFormation console](#).

6.4.3.17. Creating the Ingress DNS Records

If you removed the DNS Zone configuration, manually create DNS records that point to the Ingress load balancer. You can create either a wildcard record or specific records. While the following procedure uses A records, you can use other record types that you require, such as CNAME or alias.

Prerequisites

- You deployed an OpenShift Container Platform cluster on Amazon Web Services (AWS) that uses infrastructure that you provisioned.
- You installed the OpenShift CLI (**oc**).
- You installed the **jq** package.

- You downloaded the AWS CLI and installed it on your computer. See [Install the AWS CLI Using the Bundled Installer \(Linux, macOS, or Unix\)](#).

Procedure

- Determine the routes to create.

- To create a wildcard record, use `*.apps.<cluster_name>.<domain_name>`, where `<cluster_name>` is your cluster name, and `<domain_name>` is the Route 53 base domain for your OpenShift Container Platform cluster.
- To create specific records, you must create a record for each route that your cluster uses, as shown in the output of the following command:

```
$ oc get --all-namespaces -o jsonpath='{range .items[*]}{range .status.ingress[*]}{.host} {"\n"}}{end}{end}' routes
```

Example output

```
oauth-openshift.apps.<cluster_name>.<domain_name>
console-openshift-console.apps.<cluster_name>.<domain_name>
downloads-openshift-console.apps.<cluster_name>.<domain_name>
alertmanager-main-openshift-monitoring.apps.<cluster_name>.<domain_name>
prometheus-k8s-openshift-monitoring.apps.<cluster_name>.<domain_name>
```

- Retrieve the Ingress Operator load balancer status and note the value of the external IP address that it uses, which is shown in the **EXTERNAL-IP** column:

```
$ oc -n openshift-ingress get service router-default
```

Example output

NAME	TYPE	CLUSTER-IP	EXTERNAL-IP	PORT(S)
router-default	LoadBalancer	172.30.62.215	ab3...28.us-east-2.elb.amazonaws.com	80:31499/TCP,443:30693/TCP
		5m		

- Locate the hosted zone ID for the load balancer:

```
$ aws elb describe-load-balancers | jq -r '.LoadBalancerDescriptions[] | select(.DNSName == "<external_ip>").CanonicalHostedZoneNameID' 1
```

- For `<external_ip>`, specify the value of the external IP address of the Ingress Operator load balancer that you obtained.

Example output

```
Z3AADJGX6KTTL2
```

The output of this command is the load balancer hosted zone ID.

- Obtain the public hosted zone ID for your cluster's domain:

-

```
$ aws route53 list-hosted-zones-by-name \
  --dns-name "<domain_name>" 1
  --query 'HostedZones[? Config.PrivateZone != `true` && Name ==
`<domain_name>.`].Id' 2
  --output text
```

- 1** **2** For **<domain_name>**, specify the Route 53 base domain for your OpenShift Container Platform cluster.

Example output

```
/hostedzone/Z3URY6TWQ91KVV
```

The public hosted zone ID for your domain is shown in the command output. In this example, it is **Z3URY6TWQ91KVV**.

5. Add the alias records to your private zone:

```
$ aws route53 change-resource-record-sets --hosted-zone-id "<private_hosted_zone_id>" --
change-batch '{ 1
> "Changes": [
> {
>   "Action": "CREATE",
>   "ResourceRecordSet": {
>     "Name": "\\052.apps.<cluster_domain>", 2
>     "Type": "A",
>     "AliasTarget": {
>       "HostedZoneId": "<hosted_zone_id>", 3
>       "DNSName": "<external_ip>.", 4
>       "EvaluateTargetHealth": false
>     }
>   }
> }
> ]
> }'
```

- 1** For **<private_hosted_zone_id>**, specify the value from the output of the CloudFormation template for DNS and load balancing.
- 2** For **<cluster_domain>**, specify the domain or subdomain that you use with your OpenShift Container Platform cluster.
- 3** For **<hosted_zone_id>**, specify the public hosted zone ID for the load balancer that you obtained.
- 4** For **<external_ip>**, specify the value of the external IP address of the Ingress Operator load balancer. Ensure that you include the trailing period (.) in this parameter value.

6. Add the records to your public zone:

```
$ aws route53 change-resource-record-sets --hosted-zone-id "<public_hosted_zone_id>" --
change-batch '{ 1
```

```

> "Changes": [
> {
>   "Action": "CREATE",
>   "ResourceRecordSet": {
>     "Name": "\\052.apps.<cluster_domain>", 2
>     "Type": "A",
>     "AliasTarget":{
>       "HostedZoneId": "<hosted_zone_id>", 3
>       "DNSName": "<external_ip>.", 4
>       "EvaluateTargetHealth": false
>     }
>   }
> }
> ]
> }'

```

- 1 For **<public_hosted_zone_id>**, specify the public hosted zone for your domain.
- 2 For **<cluster_domain>**, specify the domain or subdomain that you use with your OpenShift Container Platform cluster.
- 3 For **<hosted_zone_id>**, specify the public hosted zone ID for the load balancer that you obtained.
- 4 For **<external_ip>**, specify the value of the external IP address of the Ingress Operator load balancer. Ensure that you include the trailing period (.) in this parameter value.

6.4.3.18. Completing an AWS installation on user-provisioned infrastructure

After you start the OpenShift Container Platform installation on Amazon Web Service (AWS) user-provisioned infrastructure, monitor the deployment to completion.

Prerequisites

- You removed the bootstrap node for an OpenShift Container Platform cluster on user-provisioned AWS infrastructure.
- You installed the **oc** CLI.

Procedure

- From the directory that contains the installation program, complete the cluster installation:

```
$ ./openshift-install --dir <installation_directory> wait-for install-complete 1
```

- 1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

Example output

```

INFO Waiting up to 40m0s for the cluster at https://api.mycluster.example.com:6443 to
initialize...
INFO Waiting up to 10m0s for the openshift-console route to be created...


```



```

INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 1s

```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

6.4.3.19. Logging in to the cluster by using the web console

The **kubeadmin** user exists by default after an OpenShift Container Platform installation. You can log in to your cluster as the **kubeadmin** user by using the OpenShift Container Platform web console.

Prerequisites

- You have access to the installation host.
- You completed a cluster installation and all cluster Operators are available.

Procedure

1. Obtain the password for the **kubeadmin** user from the **kubeadmin-password** file on the installation host:

```
$ cat <installation_directory>/auth/kubeadmin-password
```




NOTE

Alternatively, you can obtain the **kubeadmin** password from the **<installation_directory>/openshift_install.log** log file on the installation host.

2. List the OpenShift Container Platform web console route:

```
$ oc get routes -n openshift-console | grep 'console-openshift'
```

**NOTE**

Alternatively, you can obtain the OpenShift Container Platform route from the `<installation_directory>/openshift_install.log` log file on the installation host.

Example output

```
console console-openshift-console.apps.<cluster_name>.<base_domain> console
https reencrypt/Redirect None
```

3. Navigate to the route detailed in the output of the preceding command in a web browser and log in as the **kubeadmin** user.

Additional resources

- See [Accessing the web console](#) for more details about accessing and understanding the OpenShift Container Platform web console.

6.4.3.20. Additional resources

- See [Working with stacks](#) in the AWS documentation for more information about AWS CloudFormation stacks.

6.4.3.21. Next steps

- [Validating an installation.](#)
- [Customize your cluster.](#)
- If necessary, you can [opt out of remote health reporting](#) .
- If necessary, you can [remove cloud provider credentials](#).

6.4.4. Installing a cluster on AWS in a restricted network with user-provisioned infrastructure

In OpenShift Container Platform version 4.16, you can install a cluster on Amazon Web Services (AWS) using infrastructure that you provide and an internal mirror of the installation release content.

**IMPORTANT**

While you can install an OpenShift Container Platform cluster by using mirrored installation release content, your cluster still requires internet access to use the AWS APIs.

One way to create this infrastructure is to use the provided CloudFormation templates. You can modify the templates to customize your infrastructure or use the information that they contain to create AWS objects according to your company's policies.




IMPORTANT

The steps for performing a user-provisioned infrastructure installation are provided as an example only. Installing a cluster with infrastructure you provide requires knowledge of the cloud provider and the installation process of OpenShift Container Platform. Several CloudFormation templates are provided to assist in completing these steps or to help model your own. You are also free to create the required resources through other methods; the templates are just an example.

6.4.4.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [created a mirror registry on your mirror host](#) and obtained the **imageContentSources** data for your version of OpenShift Container Platform.



IMPORTANT

Because the installation media is on the mirror host, you can use that computer to complete all installation steps.

- You [configured an AWS account](#) to host the cluster.



IMPORTANT

If you have an AWS profile stored on your computer, it must not use a temporary session token that you generated while using a multi-factor authentication device. The cluster continues to use your current AWS credentials to create AWS resources for the entire life of the cluster, so you must use key-based, long-term credentials. To generate appropriate keys, see [Managing Access Keys for IAM Users](#) in the AWS documentation. You can supply the keys when you run the installation program.

- You [prepared the user-provisioned infrastructure](#).
- You downloaded the AWS CLI and installed it on your computer. See [Install the AWS CLI Using the Bundled Installer \(Linux, macOS, or UNIX\)](#) in the AWS documentation.
- If you use a firewall and plan to use the Telemetry service, you [configured the firewall to allow the sites](#) that your cluster requires access to.



NOTE

Be sure to also review this site list if you are configuring a proxy.

- If the cloud identity and access management (IAM) APIs are not accessible in your environment, or if you do not want to store an administrator-level credential secret in the **kube-system** namespace, you can [manually create and maintain long-term credentials](#).

6.4.4.2. About installations in restricted networks

In OpenShift Container Platform 4.16, you can perform an installation that does not require an active connection to the internet to obtain software components. Restricted network installations can be completed using installer-provisioned infrastructure or user-provisioned infrastructure, depending on the cloud platform to which you are installing the cluster.

If you choose to perform a restricted network installation on a cloud platform, you still require access to its cloud APIs. Some cloud functions, like Amazon Web Service's Route 53 DNS and IAM services, require internet access. Depending on your network, you might require less internet access for an installation on bare metal hardware, Nutanix, or on VMware vSphere.

To complete a restricted network installation, you must create a registry that mirrors the contents of the OpenShift image registry and contains the installation media. You can create this registry on a mirror host, which can access both the internet and your closed network, or by using other methods that meet your restrictions.



IMPORTANT

Because of the complexity of the configuration for user-provisioned installations, consider completing a standard user-provisioned infrastructure installation before you attempt a restricted network installation using user-provisioned infrastructure. Completing this test installation might make it easier to isolate and troubleshoot any issues that might arise during your installation in a restricted network.

6.4.4.2.1. Additional limits

Clusters in restricted networks have the following additional limitations and restrictions:

- The **ClusterVersion** status includes an **Unable to retrieve available updates** error.
- By default, you cannot use the contents of the Developer Catalog because you cannot access the required image stream tags.

6.4.4.3. Creating the installation files for AWS

To install OpenShift Container Platform on Amazon Web Services (AWS) using user-provisioned infrastructure, you must generate the files that the installation program needs to deploy your cluster and modify them so that the cluster creates only the machines that it will use. You generate and customize the **install-config.yaml** file, Kubernetes manifests, and Ignition config files. You also have the option to first set up a separate **var** partition during the preparation phases of installation.

6.4.4.3.1. Optional: Creating a separate **/var** partition

It is recommended that disk partitioning for OpenShift Container Platform be left to the installer. However, there are cases where you might want to create separate partitions in a part of the filesystem that you expect to grow.


OpenShift Container Platform supports the addition of a single partition to attach storage to either the **/var** partition or a subdirectory of **/var**. For example:

- **/var/lib/containers**: Holds container-related content that can grow as more images and containers are added to a system.
- **/var/lib/etcd**: Holds data that you might want to keep separate for purposes such as performance optimization of etcd storage.

- **/var**: Holds data that you might want to keep separate for purposes such as auditing.

Storing the contents of a **/var** directory separately makes it easier to grow storage for those areas as needed and reinstall OpenShift Container Platform at a later date and keep that data intact. With this method, you will not have to pull all your containers again, nor will you have to copy massive log files when you update systems.

Because **/var** must be in place before a fresh installation of Red Hat Enterprise Linux CoreOS (RHCOS), the following procedure sets up the separate **/var** partition by creating a machine config manifest that is inserted during the **openshift-install** preparation phases of an OpenShift Container Platform installation.



IMPORTANT

If you follow the steps to create a separate **/var** partition in this procedure, it is not necessary to create the Kubernetes manifest and Ignition config files again as described later in this section.

Procedure

1. Create a directory to hold the OpenShift Container Platform installation files:

```
$ mkdir $HOME/clusterconfig
```

2. Run **openshift-install** to create a set of files in the **manifest** and **openshift** subdirectories. Answer the system questions as you are prompted:

```
$ openshift-install create manifests --dir $HOME/clusterconfig
```

Example output

```
? SSH Public Key ...
INFO Credentials loaded from the "myprofile" profile in file "/home/myuser/.aws/credentials"
INFO Consuming Install Config from target directory
INFO Manifests created in: $HOME/clusterconfig/manifests and
$HOME/clusterconfig/openshift
```

3. Optional: Confirm that the installation program created manifests in the **clusterconfig/openshift** directory:

```
$ ls $HOME/clusterconfig/openshift/
```

Example output

```
99_kubeadmin-password-secret.yaml
99_openshift-cluster-api_master-machines-0.yaml
99_openshift-cluster-api_master-machines-1.yaml
99_openshift-cluster-api_master-machines-2.yaml
...
```

4. Create a Butane config that configures the additional partition. For example, name the file **\$HOME/clusterconfig/98-var-partition.bu**, change the disk device name to the name of the storage device on the **worker** systems, and set the storage size as appropriate. This example

places the **/var** directory on a separate partition:

```
variant: openshift
version: 4.16.0
metadata:
  labels:
    machineconfiguration.openshift.io/role: worker
  name: 98-var-partition
storage:
  disks:
    - device: /dev/disk/by-id/<device_name> ❶
      partitions:
        - label: var
          start_mib: <partition_start_offset> ❷
          size_mib: <partition_size> ❸
          number: 5
      filesystems:
        - device: /dev/disk/by-partlabel/var
          path: /var
          format: xfs
          mount_options: [defaults, prjquota] ❹
          with_mount_unit: true
```

- ❶ The storage device name of the disk that you want to partition.
- ❷ When adding a data partition to the boot disk, a minimum value of 25000 MiB (Mebibytes) is recommended. The root file system is automatically resized to fill all available space up to the specified offset. If no value is specified, or if the specified value is smaller than the recommended minimum, the resulting root file system will be too small, and future reinstalls of RHCOS might overwrite the beginning of the data partition.
- ❸ The size of the data partition in mebibytes.
- ❹ The **prjquota** mount option must be enabled for filesystems used for container storage.



NOTE

When creating a separate **/var** partition, you cannot use different instance types for worker nodes, if the different instance types do not have the same device name.

5. Create a manifest from the Butane config and save it to the **clusterconfig/openshift** directory. For example, run the following command:

```
$ butane $HOME/clusterconfig/98-var-partition.bu -o $HOME/clusterconfig/openshift/98-var-partition.yaml
```

6. Run **openshift-install** again to create Ignition configs from a set of files in the **manifest** and **openshift** subdirectories:

```
$ openshift-install create ignition-configs --dir $HOME/clusterconfig
$ ls $HOME/clusterconfig/
auth bootstrap.ign master.ign metadata.json worker.ign
```

Now you can use the Ignition config files as input to the installation procedures to install Red Hat Enterprise Linux CoreOS (RHCOS) systems.

6.4.4.3.2. Creating the installation configuration file

Generate and customize the installation configuration file that the installation program needs to deploy your cluster.

Prerequisites

- You obtained the OpenShift Container Platform installation program for user-provisioned infrastructure and the pull secret for your cluster. For a restricted network installation, these files are on your mirror host.
- You checked that you are deploying your cluster to an AWS Region with an accompanying Red Hat Enterprise Linux CoreOS (RHCOS) AMI published by Red Hat. If you are deploying to an AWS Region that requires a custom AMI, such as an AWS GovCloud Region, you must create the **install-config.yaml** file manually.

Procedure

1. Create the **install-config.yaml** file.
 - a. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install create install-config --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the directory name to store the files that the installation program creates.



IMPORTANT

Specify an empty directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

- b. At the prompts, provide the configuration details for your cloud:
 - i. Optional: Select an SSH key to use to access your cluster machines.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- ii. Select **aws** as the platform to target.

- iii. If you do not have an AWS profile stored on your computer, enter the AWS access key ID and secret access key for the user that you configured to run the installation program.



NOTE

The AWS access key ID and secret access key are stored in `~/.aws/credentials` in the home directory of the current user on the installation host. You are prompted for the credentials by the installation program if the credentials for the exported profile are not present in the file. Any credentials that you provide to the installation program are stored in the file.

- iv. Select the AWS Region to deploy the cluster to.
 - v. Select the base domain for the Route 53 service that you configured for your cluster.
 - vi. Enter a descriptive name for your cluster.
 - vii. Paste the [pull secret from Red Hat OpenShift Cluster Manager](#) .
2. Edit the `install-config.yaml` file to give the additional information that is required for an installation in a restricted network.

- a. Update the `pullSecret` value to contain the authentication information for your registry:

```
pullSecret: '{"auths":{"<local_registry>": {"auth": "<credentials>","email":
"you@example.com"}}}'
```

For `<local_registry>`, specify the registry domain name, and optionally the port, that your mirror registry uses to serve content. For example `registry.example.com` or `registry.example.com:5000`. For `<credentials>`, specify the base64-encoded user name and password for your mirror registry.

- b. Add the `additionalTrustBundle` parameter and value. The value must be the contents of the certificate file that you used for your mirror registry. The certificate file can be an existing, trusted certificate authority or the self-signed certificate that you generated for the mirror registry.

```
additionalTrustBundle: |
  -----BEGIN CERTIFICATE-----
  /-----END CERTIFICATE-----
```

- c. Add the image content resources:

```
imageContentSources:
- mirrors:
  - <local_registry>/<local_repository_name>/release
  source: quay.io/openshift-release-dev/ocp-release
- mirrors:
  - <local_registry>/<local_repository_name>/release
  source: quay.io/openshift-release-dev/ocp-v4.0-art-dev
```


Use the **imageContentSources** section from the output of the command to mirror the repository or the values that you used when you mirrored the content from the media that you brought into your restricted network.

- d. Optional: Set the publishing strategy to **Internal**:

```
publish: Internal
```

By setting this option, you create an internal Ingress Controller and a private load balancer.

3. Optional: Back up the **install-config.yaml** file.



IMPORTANT

The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

Additional resources

- See [Configuration and credential file settings](#) in the AWS documentation for more information about AWS profile and credential configuration.

6.4.4.3.3. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
```

```

proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: ec2.<aws_region>.amazonaws.com,elasticloadbalancing.
<aws_region>.amazonaws.com,s3.<aws_region>.amazonaws.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5

```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations. If you have added the Amazon **EC2,Elastic Load Balancing**, and **S3** VPC endpoints to your VPC, you must add these endpoints to the **noProxy** field.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.



NOTE

The installation program does not support the proxy **readinessEndpoints** field.



NOTE

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.



NOTE

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

6.4.4.3.4. Creating the Kubernetes manifest and Ignition config files

Because you must modify some cluster definition files and manually start the cluster machines, you must generate the Kubernetes manifest and Ignition config files that the cluster needs to configure the machines.

The installation configuration file transforms into the Kubernetes manifests. The manifests wrap into the Ignition configuration files, which are later used to configure the cluster machines.



IMPORTANT

- The Ignition config files that the OpenShift Container Platform installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

Prerequisites

- You obtained the OpenShift Container Platform installation program. For a restricted network installation, these files are on your mirror host.
- You created the **install-config.yaml** installation configuration file.

Procedure

1. Change to the directory that contains the OpenShift Container Platform installation program and generate the Kubernetes manifests for the cluster:

```
$ ./openshift-install create manifests --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the installation directory that contains the **install-config.yaml** file you created.

2. Remove the Kubernetes manifest files that define the control plane machines:

```
$ rm -f <installation_directory>/openshift/99_openshift-cluster-api_master-machines-*.yaml
```

By removing these files, you prevent the cluster from automatically generating control plane machines.

- Remove the Kubernetes manifest files that define the control plane machine set:

```
$ rm -f <installation_directory>/openshift/99_openshift-machine-api_master-control-plane-machine-set.yaml
```

```
$ rm -f <installation_directory>/openshift/99_openshift-cluster-api_worker-machineset-*.yaml
```



IMPORTANT

If you disabled the **MachineAPI** capability when installing a cluster on user-provisioned infrastructure, you must remove the Kubernetes manifest files that define the worker machines. Otherwise, your cluster fails to install.

Because you create and manage the worker machines yourself, you do not need to initialize these machines.

- Check that the **mastersSchedulable** parameter in the `<installation_directory>/manifests/cluster-scheduler-02-config.yml` Kubernetes manifest file is set to **false**. This setting prevents pods from being scheduled on the control plane machines:
 - Open the `<installation_directory>/manifests/cluster-scheduler-02-config.yml` file.
 - Locate the **mastersSchedulable** parameter and ensure that it is set to **false**.
 - Save and exit the file.
- Optional: If you do not want [the Ingress Operator](#) to create DNS records on your behalf, remove the **privateZone** and **publicZone** sections from the `<installation_directory>/manifests/cluster-dns-02-config.yml` DNS configuration file:

```
apiVersion: config.openshift.io/v1
kind: DNS
metadata:
  creationTimestamp: null
  name: cluster
spec:
  baseDomain: example.openshift.com
  privateZone: 1
    id: mycluster-100419-private-zone
  publicZone: 2
    id: example.openshift.com
status: {}
```

- 1 2 Remove this section completely.

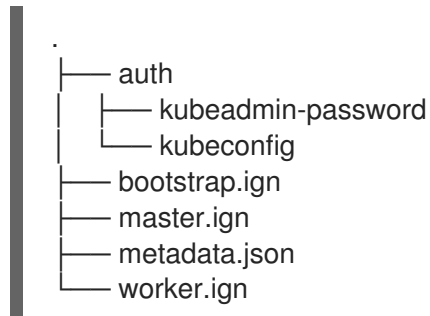
If you do so, you must add ingress DNS records manually in a later step.

- To create the Ignition configuration files, run the following command from the directory that contains the installation program:

```
$ ./openshift-install create ignition-configs --dir <installation_directory> 1
```

- 1 For `<installation_directory>`, specify the same installation directory.

Ignition config files are created for the bootstrap, control plane, and compute nodes in the installation directory. The **kubeadmin-password** and **kubeconfig** files are created in the `./<installation_directory>/auth` directory:



Additional resources

- [Manually creating long-term credentials](#)

6.4.4.4. Extracting the infrastructure name

The Ignition config files contain a unique cluster identifier that you can use to uniquely identify your cluster in Amazon Web Services (AWS). The infrastructure name is also used to locate the appropriate AWS resources during an OpenShift Container Platform installation. The provided CloudFormation templates contain references to this infrastructure name, so you must extract it.

Prerequisites

- You obtained the OpenShift Container Platform installation program and the pull secret for your cluster.
- You generated the Ignition config files for your cluster.
- You installed the **jq** package.

Procedure

- To extract and view the infrastructure name from the Ignition config file metadata, run the following command:

```
$ jq -r .infraID <installation_directory>/metadata.json 1
```

- 1 For `<installation_directory>`, specify the path to the directory that you stored the installation files in.

Example output

```
openshift-vw9j6 1
```

- 1 The output of this command is your cluster name and a random string.

6.4.4.5. Creating a VPC in AWS

You must create a Virtual Private Cloud (VPC) in Amazon Web Services (AWS) for your OpenShift Container Platform cluster to use. You can customize the VPC to meet your requirements, including VPN and route tables.

You can use the provided CloudFormation template and a custom parameter file to create a stack of AWS resources that represent the VPC.



NOTE

If you do not use the provided CloudFormation template to create your AWS infrastructure, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- You configured an AWS account.
- You added your AWS keys and region to your local AWS profile by running **aws configure**.
- You generated the Ignition config files for your cluster.

Procedure

1. Create a JSON file that contains the parameter values that the template requires:

```
[
  {
    "ParameterKey": "VpcCidr", 1
    "ParameterValue": "10.0.0.0/16" 2
  },
  {
    "ParameterKey": "AvailabilityZoneCount", 3
    "ParameterValue": "1" 4
  },
  {
    "ParameterKey": "SubnetBits", 5
    "ParameterValue": "12" 6
  }
]
```

- 1** The CIDR block for the VPC.
- 2** Specify a CIDR block in the format **x.x.x.x/16-24**.
- 3** The number of availability zones to deploy the VPC in.
- 4** Specify an integer between **1** and **3**.
- 5** The size of each subnet in each availability zone.
- 6** Specify an integer between **5** and **13**, where **5** is **/27** and **13** is **/19**.

- Copy the template from the **CloudFormation template for the VPC** section of this topic and save it as a YAML file on your computer. This template describes the VPC that your cluster requires.
- Launch the CloudFormation template to create a stack of AWS resources that represent the VPC:



IMPORTANT

You must enter the command on a single line.

```
$ aws cloudformation create-stack --stack-name <name> 1
  --template-body file://<template>.yaml 2
  --parameters file://<parameters>.json 3
```

- <name>** is the name for the CloudFormation stack, such as **cluster-vpc**. You need the name of this stack if you remove the cluster.
- <template>** is the relative path to and name of the CloudFormation template YAML file that you saved.
- <parameters>** is the relative path to and name of the CloudFormation parameters JSON file.

Example output

```
arn:aws:cloudformation:us-east-1:269333783861:stack/cluster-vpc/dbedae40-2fd3-11eb-820e-12a48460849f
```

- Confirm that the template components exist:

```
$ aws cloudformation describe-stacks --stack-name <name>
```

After the **StackStatus** displays **CREATE_COMPLETE**, the output displays values for the following parameters. You must provide these parameter values to the other CloudFormation templates that you run to create your cluster:

VpcId	The ID of your VPC.
PublicSubnetIds	The IDs of the new public subnets.
PrivateSubnetIds	The IDs of the new private subnets.

6.4.4.5.1. CloudFormation template for the VPC

You can use the following CloudFormation template to deploy the VPC that you need for your OpenShift Container Platform cluster.

Example 6.78. CloudFormation template for the VPC

```

AWSTemplateFormatVersion: 2010-09-09
Description: Template for Best Practice VPC with 1-3 AZs

Parameters:
  VpcCidr:
    AllowedPattern: ^(((0-9){1,3}(0-9){0,2}|2[0-4][0-9]{2}|25[0-5]).){3}((0-9){1,3}(0-9){0,2}|2[0-4][0-9]{2}|25[0-5])(\.(1[6-9]|2[0-4]))$
    ConstraintDescription: CIDR block parameter must be in the form x.x.x.x/16-24.
    Default: 10.0.0.0/16
    Description: CIDR block for VPC.
    Type: String
  AvailabilityZoneCount:
    ConstraintDescription: "The number of availability zones. (Min: 1, Max: 3)"
    MinValue: 1
    MaxValue: 3
    Default: 1
    Description: "How many AZs to create VPC subnets for. (Min: 1, Max: 3)"
    Type: Number
  SubnetBits:
    ConstraintDescription: CIDR block parameter must be in the form x.x.x.x/19-27.
    MinValue: 5
    MaxValue: 13
    Default: 12
    Description: "Size of each subnet to create within the availability zones. (Min: 5 = /27, Max: 13 = /19)"
    Type: Number

Metadata:
  AWS::CloudFormation::Interface:
    ParameterGroups:
      - Label:
          default: "Network Configuration"
        Parameters:
          - VpcCidr
          - SubnetBits
      - Label:
          default: "Availability Zones"
        Parameters:
          - AvailabilityZoneCount
    ParameterLabels:
      AvailabilityZoneCount:
        default: "Availability Zone Count"
      VpcCidr:
        default: "VPC CIDR"
      SubnetBits:
        default: "Bits Per Subnet"

Conditions:
  DoAz3: !Equals [3, !Ref AvailabilityZoneCount]
  DoAz2: !Or [!Equals [2, !Ref AvailabilityZoneCount], Condition: DoAz3]

Resources:
  VPC:
    Type: "AWS::EC2::VPC"

```



```

Properties:
  EnableDnsSupport: "true"
  EnableDnsHostnames: "true"
  CidrBlock: !Ref VpcCidr
PublicSubnet:
  Type: "AWS::EC2::Subnet"
  Properties:
    Vpclid: !Ref VPC
    CidrBlock: !Select [0, !Cidr [!Ref VpcCidr, 6, !Ref SubnetBits]]
    AvailabilityZone: !Select
      - 0
      - Fn::GetAZs: !Ref "AWS::Region"
PublicSubnet2:
  Type: "AWS::EC2::Subnet"
  Condition: DoAz2
  Properties:
    Vpclid: !Ref VPC
    CidrBlock: !Select [1, !Cidr [!Ref VpcCidr, 6, !Ref SubnetBits]]
    AvailabilityZone: !Select
      - 1
      - Fn::GetAZs: !Ref "AWS::Region"
PublicSubnet3:
  Type: "AWS::EC2::Subnet"
  Condition: DoAz3
  Properties:
    Vpclid: !Ref VPC
    CidrBlock: !Select [2, !Cidr [!Ref VpcCidr, 6, !Ref SubnetBits]]
    AvailabilityZone: !Select
      - 2
      - Fn::GetAZs: !Ref "AWS::Region"
InternetGateway:
  Type: "AWS::EC2::InternetGateway"
GatewayToInternet:
  Type: "AWS::EC2::VPCGatewayAttachment"
  Properties:
    Vpclid: !Ref VPC
    InternetGatewayId: !Ref InternetGateway
PublicRouteTable:
  Type: "AWS::EC2::RouteTable"
  Properties:
    Vpclid: !Ref VPC
PublicRoute:
  Type: "AWS::EC2::Route"
  DependsOn: GatewayToInternet
  Properties:
    RouteTableId: !Ref PublicRouteTable
    DestinationCidrBlock: 0.0.0.0/0
    GatewayId: !Ref InternetGateway
PublicSubnetRouteTableAssociation:
  Type: "AWS::EC2::SubnetRouteTableAssociation"
  Properties:
    SubnetId: !Ref PublicSubnet
    RouteTableId: !Ref PublicRouteTable
PublicSubnetRouteTableAssociation2:
  Type: "AWS::EC2::SubnetRouteTableAssociation"
  Condition: DoAz2

```

```

Properties:
  SubnetId: !Ref PublicSubnet2
  RouteTableId: !Ref PublicRouteTable
PublicSubnetRouteTableAssociation3:
  Condition: DoAz3
  Type: "AWS::EC2::SubnetRouteTableAssociation"
Properties:
  SubnetId: !Ref PublicSubnet3
  RouteTableId: !Ref PublicRouteTable
PrivateSubnet:
  Type: "AWS::EC2::Subnet"
Properties:
  VpcId: !Ref VPC
  CidrBlock: !Select [3, !Cidr [!Ref VpcCidr, 6, !Ref SubnetBits]]
  AvailabilityZone: !Select
  - 0
  - Fn::GetAZs: !Ref "AWS::Region"
PrivateRouteTable:
  Type: "AWS::EC2::RouteTable"
Properties:
  VpcId: !Ref VPC
PrivateSubnetRouteTableAssociation:
  Type: "AWS::EC2::SubnetRouteTableAssociation"
Properties:
  SubnetId: !Ref PrivateSubnet
  RouteTableId: !Ref PrivateRouteTable
NAT:
  DependsOn:
  - GatewayToInternet
  Type: "AWS::EC2::NatGateway"
Properties:
  AllocationId:
    "Fn::GetAtt":
    - EIP
    - AllocationId
  SubnetId: !Ref PublicSubnet
EIP:
  Type: "AWS::EC2::EIP"
Properties:
  Domain: vpc
Route:
  Type: "AWS::EC2::Route"
Properties:
  RouteTableId:
    Ref: PrivateRouteTable
  DestinationCidrBlock: 0.0.0.0/0
  NatGatewayId:
    Ref: NAT
PrivateSubnet2:
  Type: "AWS::EC2::Subnet"
  Condition: DoAz2
Properties:
  VpcId: !Ref VPC
  CidrBlock: !Select [4, !Cidr [!Ref VpcCidr, 6, !Ref SubnetBits]]
  AvailabilityZone: !Select
  - 1

```

```

- Fn::GetAZs: !Ref "AWS::Region"
PrivateRouteTable2:
  Type: "AWS::EC2::RouteTable"
  Condition: DoAz2
  Properties:
    VpId: !Ref VPC
PrivateSubnetRouteTableAssociation2:
  Type: "AWS::EC2::SubnetRouteTableAssociation"
  Condition: DoAz2
  Properties:
    SubnetId: !Ref PrivateSubnet2
    RouteTableId: !Ref PrivateRouteTable2
NAT2:
  DependsOn:
  - GatewayToInternet
  Type: "AWS::EC2::NatGateway"
  Condition: DoAz2
  Properties:
    AllocationId:
      "Fn::GetAtt":
      - EIP2
      - AllocationId
    SubnetId: !Ref PublicSubnet2
EIP2:
  Type: "AWS::EC2::EIP"
  Condition: DoAz2
  Properties:
    Domain: vpc
Route2:
  Type: "AWS::EC2::Route"
  Condition: DoAz2
  Properties:
    RouteTableId:
      Ref: PrivateRouteTable2
    DestinationCidrBlock: 0.0.0.0/0
    NatGatewayId:
      Ref: NAT2
PrivateSubnet3:
  Type: "AWS::EC2::Subnet"
  Condition: DoAz3
  Properties:
    VpId: !Ref VPC
    CidrBlock: !Select [5, !Cidr [!Ref VpcCidr, 6, !Ref SubnetBits]]
    AvailabilityZone: !Select
      - 2
      - Fn::GetAZs: !Ref "AWS::Region"
PrivateRouteTable3:
  Type: "AWS::EC2::RouteTable"
  Condition: DoAz3
  Properties:
    VpId: !Ref VPC
PrivateSubnetRouteTableAssociation3:
  Type: "AWS::EC2::SubnetRouteTableAssociation"
  Condition: DoAz3
  Properties:
    SubnetId: !Ref PrivateSubnet3

```

```

RouteTableId: !Ref PrivateRouteTable3
NAT3:
  DependsOn:
  - GatewayToInternet
  Type: "AWS::EC2::NatGateway"
  Condition: DoAz3
  Properties:
    AllocationId:
      "Fn::GetAtt":
      - EIP3
      - AllocationId
    SubnetId: !Ref PublicSubnet3
EIP3:
  Type: "AWS::EC2::EIP"
  Condition: DoAz3
  Properties:
    Domain: vpc
Route3:
  Type: "AWS::EC2::Route"
  Condition: DoAz3
  Properties:
    RouteTableId:
      Ref: PrivateRouteTable3
    DestinationCidrBlock: 0.0.0.0/0
    NatGatewayId:
      Ref: NAT3
S3Endpoint:
  Type: AWS::EC2::VPCEndpoint
  Properties:
    PolicyDocument:
      Version: 2012-10-17
      Statement:
      - Effect: Allow
        Principal: '*'
        Action:
          - '*'
        Resource:
          - '*'
    RouteTableIds:
      - !Ref PublicRouteTable
      - !Ref PrivateRouteTable
      - !If [DoAz2, !Ref PrivateRouteTable2, !Ref "AWS::NoValue"]
      - !If [DoAz3, !Ref PrivateRouteTable3, !Ref "AWS::NoValue"]
    ServiceName: !Join
      - "
      - - com.amazonaws.
      - - !Ref 'AWS::Region'
      - - .s3
    VpcId: !Ref VPC

Outputs:
VpcId:
  Description: ID of the new VPC.
  Value: !Ref VPC
PublicSubnetIds:
  Description: Subnet IDs of the public subnets.

```

```

Value:
  !Join [
    ",",
    [!Ref PublicSubnet, !If [DoAz2, !Ref PublicSubnet2, !Ref "AWS::NoValue"], !If [DoAz3, !Ref
PublicSubnet3, !Ref "AWS::NoValue"]]
  ]
PrivateSubnetIds:
  Description: Subnet IDs of the private subnets.
  Value:
    !Join [
      ",",
      [!Ref PrivateSubnet, !If [DoAz2, !Ref PrivateSubnet2, !Ref "AWS::NoValue"], !If [DoAz3, !Ref
PrivateSubnet3, !Ref "AWS::NoValue"]]
    ]
PublicRouteTableId:
  Description: Public Route table ID
  Value: !Ref PublicRouteTable
PrivateRouteTableIds:
  Description: Private Route table IDs
  Value:
    !Join [
      ",",
      [
        !Join ["=", [
          !Select [0, "Fn::GetAZs": !Ref "AWS::Region"],
          !Ref PrivateRouteTable
        ]],
        !If [DoAz2,
          !Join ["=", [!Select [1, "Fn::GetAZs": !Ref "AWS::Region"], !Ref PrivateRouteTable2]],
          !Ref "AWS::NoValue"
        ],
        !If [DoAz3,
          !Join ["=", [!Select [2, "Fn::GetAZs": !Ref "AWS::Region"], !Ref PrivateRouteTable3]],
          !Ref "AWS::NoValue"
        ]
      ]
    ]
  ]
]

```

6.4.4.6. Creating networking and load balancing components in AWS

You must configure networking and classic or network load balancing in Amazon Web Services (AWS) that your OpenShift Container Platform cluster can use.

You can use the provided CloudFormation template and a custom parameter file to create a stack of AWS resources. The stack represents the networking and load balancing components that your OpenShift Container Platform cluster requires. The template also creates a hosted zone and subnet tags.

You can run the template multiple times within a single Virtual Private Cloud (VPC).



NOTE

If you do not use the provided CloudFormation template to create your AWS infrastructure, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- You configured an AWS account.
- You added your AWS keys and region to your local AWS profile by running **aws configure**.
- You generated the Ignition config files for your cluster.
- You created and configured a VPC and associated subnets in AWS.

Procedure

1. Obtain the hosted zone ID for the Route 53 base domain that you specified in the **install-config.yaml** file for your cluster. You can obtain details about your hosted zone by running the following command:

```
$ aws route53 list-hosted-zones-by-name --dns-name <route53_domain> 1
```

- 1 For the **<route53_domain>**, specify the Route 53 base domain that you used when you generated the **install-config.yaml** file for the cluster.

Example output

```
mycluster.example.com. False 100
HOSTEDZONES 65F8F38E-2268-B835-E15C-AB55336FCBFA
/hostedzone/Z21IXYZABCZ2A4 mycluster.example.com. 10
```

In the example output, the hosted zone ID is **Z21IXYZABCZ2A4**.

2. Create a JSON file that contains the parameter values that the template requires:


```
[
  {
    "ParameterKey": "ClusterName", 1
    "ParameterValue": "mycluster" 2
  },
  {
    "ParameterKey": "InfrastructureName", 3
    "ParameterValue": "mycluster-<random_string>" 4
  },
  {
    "ParameterKey": "HostedZoneId", 5
    "ParameterValue": "<random_string>" 6
  },
  {
    "ParameterKey": "HostedZoneName", 7
  }
]
```

```

    "ParameterValue": "example.com" 8
  },
  {
    "ParameterKey": "PublicSubnets", 9
    "ParameterValue": "subnet-<random_string>" 10
  },
  {
    "ParameterKey": "PrivateSubnets", 11
    "ParameterValue": "subnet-<random_string>" 12
  },
  {
    "ParameterKey": "VpcId", 13
    "ParameterValue": "vpc-<random_string>" 14
  }
]

```

- 1 A short, representative cluster name to use for hostnames, etc.
 - 2 Specify the cluster name that you used when you generated the **install-config.yaml** file for the cluster.
 - 3 The name for your cluster infrastructure that is encoded in your Ignition config files for the cluster.
 - 4 Specify the infrastructure name that you extracted from the Ignition config file metadata, which has the format **<cluster-name>-<random-string>**.
 - 5 The Route 53 public zone ID to register the targets with.
 - 6 Specify the Route 53 public zone ID, which as a format similar to **Z21IXYZABCZ2A4**. You can obtain this value from the AWS console.
 - 7 The Route 53 zone to register the targets with.
 - 8 Specify the Route 53 base domain that you used when you generated the **install-config.yaml** file for the cluster. Do not include the trailing period (.) that is displayed in the AWS console.
 - 9 The public subnets that you created for your VPC.
 - 10 Specify the **PublicSubnetIds** value from the output of the CloudFormation template for the VPC.
 - 11 The private subnets that you created for your VPC.
 - 12 Specify the **PrivateSubnetIds** value from the output of the CloudFormation template for the VPC.
 - 13 The VPC that you created for the cluster.
 - 14 Specify the **VpcId** value from the output of the CloudFormation template for the VPC.
3. Copy the template from the **CloudFormation template for the network and load balancers** section of this topic and save it as a YAML file on your computer. This template describes the networking and load balancing objects that your cluster requires.

**IMPORTANT**

If you are deploying your cluster to an AWS government or secret region, you must update the **InternalApiServerRecord** in the CloudFormation template to use **CNAME** records. Records of type **ALIAS** are not supported for AWS government regions.

4. Launch the CloudFormation template to create a stack of AWS resources that provide the networking and load balancing components:

**IMPORTANT**

You must enter the command on a single line.

```
$ aws cloudformation create-stack --stack-name <name> 1
  --template-body file://<template>.yaml 2
  --parameters file://<parameters>.json 3
  --capabilities CAPABILITY_NAMED_IAM 4
```

- 1** **<name>** is the name for the CloudFormation stack, such as **cluster-dns**. You need the name of this stack if you remove the cluster.
- 2** **<template>** is the relative path to and name of the CloudFormation template YAML file that you saved.
- 3** **<parameters>** is the relative path to and name of the CloudFormation parameters JSON file.
- 4** You must explicitly declare the **CAPABILITY_NAMED_IAM** capability because the provided template creates some **AWS::IAM::Role** resources.

Example output

```
arn:aws:cloudformation:us-east-1:269333783861:stack/cluster-dns/cd3e5de0-2fd4-11eb-5cf0-12be5c33a183
```

5. Confirm that the template components exist:

```
$ aws cloudformation describe-stacks --stack-name <name>
```

After the **StackStatus** displays **CREATE_COMPLETE**, the output displays values for the following parameters. You must provide these parameter values to the other CloudFormation templates that you run to create your cluster:

PrivateHostedZoneId	Hosted zone ID for the private DNS.
----------------------------	-------------------------------------

ExternalApiLoadBalancerName	Full name of the external API load balancer.
InternalApiLoadBalancerName	Full name of the internal API load balancer.
ApiServerDnsName	Full hostname of the API server.
RegisterNlbIpTargetsLambda	Lambda ARN useful to help register/deregister IP targets for these load balancers.
ExternalApiTargetGroupArn	ARN of external API target group.
InternalApiTargetGroupArn	ARN of internal API target group.
InternalServiceTargetGroupArn	ARN of internal service target group.

6.4.4.6.1. CloudFormation template for the network and load balancers

You can use the following CloudFormation template to deploy the networking objects and load balancers that you need for your OpenShift Container Platform cluster.

Example 6.79. CloudFormation template for the network and load balancers

AWSTemplateFormatVersion: [2010-09-09](#)

Description: Template for OpenShift Cluster Network Elements (Route53 & LBs)

Parameters:

ClusterName:

AllowedPattern: `^[a-zA-Z][a-zA-Z0-9\-\-]{0,26}$`

MaxLength: `27`

MinLength: `1`

ConstraintDescription: Cluster name must be alphanumeric, start with a letter, and have a maximum of `27` characters.

Description: A short, representative cluster name to use for host names and other identifying names.

Type: String

InfrastructureName:

AllowedPattern: `^[a-zA-Z][a-zA-Z0-9\-\-]{0,26}$`

MaxLength: 27

MinLength: 1

ConstraintDescription: Infrastructure name must be alphanumeric, start with a letter, and have a maximum of 27 characters.

Description: A short, unique cluster ID used to tag cloud resources and identify items owned or used by the cluster.

Type: String

HostedZoneId:

Description: The Route53 public zone ID to register the targets with, such as Z21IXYZABCZ2A4.

Type: String

HostedZoneName:

Description: The Route53 zone to register the targets with, such as example.com. Omit the trailing period.

Type: String

Default: "example.com"

PublicSubnets:

Description: The internet-facing subnets.

Type: List<AWS::EC2::Subnet::Id>

PrivateSubnets:

Description: The internal subnets.

Type: List<AWS::EC2::Subnet::Id>

VpcId:

Description: The VPC-scoped resources will belong to this VPC.

Type: AWS::EC2::VPC::Id

Metadata:

AWS::CloudFormation::Interface:

ParameterGroups:

- Label:

default: "Cluster Information"

Parameters:

- ClusterName

- InfrastructureName

- Label:

default: "Network Configuration"

Parameters:

- VpcId

- PublicSubnets

- PrivateSubnets

- Label:

default: "DNS"

Parameters:

- HostedZoneName

- HostedZoneId

ParameterLabels:

ClusterName:

default: "Cluster Name"

InfrastructureName:

default: "Infrastructure Name"

VpcId:

default: "VPC ID"

PublicSubnets:

default: "Public Subnets"

PrivateSubnets:

default: "Private Subnets"

HostedZoneName:
 default: "Public Hosted Zone Name"
 HostedZoneId:
 default: "Public Hosted Zone ID"

Resources:

ExtApiElb:

Type: AWS::ElasticLoadBalancingV2::LoadBalancer
 Properties:
 Name: !Join ["-", [!Ref InfrastructureName, "ext"]]
 IpAddressType: ipv4
 Subnets: !Ref PublicSubnets
 Type: network

IntApiElb:

Type: AWS::ElasticLoadBalancingV2::LoadBalancer
 Properties:
 Name: !Join ["-", [!Ref InfrastructureName, "int"]]
 Scheme: internal
 IpAddressType: ipv4
 Subnets: !Ref PrivateSubnets
 Type: network

IntDns:

Type: "AWS::Route53::HostedZone"
 Properties:
 HostedZoneConfig:
 Comment: "Managed by CloudFormation"
 Name: !Join [".", [!Ref ClusterName, !Ref HostedZoneName]]
 HostedZoneTags:
 - Key: Name
 Value: !Join ["-", [!Ref InfrastructureName, "int"]]
 - Key: !Join ["", ["kubernetes.io/cluster/", !Ref InfrastructureName]]
 Value: "owned"
 VPCs:
 - VPCId: !Ref Vpclid
 VPCRegion: !Ref "AWS::Region"

ExternalApiServerRecord:

Type: AWS::Route53::RecordSetGroup
 Properties:
 Comment: Alias record for the API server
 HostedZoneId: !Ref HostedZoneId
 RecordSets:
 - Name:
 !Join [
 ".",
 ["api", !Ref ClusterName, !Join ["", [!Ref HostedZoneName, "."]]],
]
 Type: A
 AliasTarget:
 HostedZoneId: !GetAtt ExtApiElb.CanonicalHostedZoneID
 DNSName: !GetAtt ExtApiElb.DNSName

InternalApiServerRecord:

Type: AWS::Route53::RecordSetGroup

Properties:

Comment: Alias record for the API server

HostedZoneId: !Ref IntDns

RecordSets:

- Name:

```
!Join [
  ".",
  ["api", !Ref ClusterName, !Join ["", [!Ref HostedZoneName, "."]]],
]
```

Type: A

AliasTarget:

HostedZoneId: !GetAtt IntApiElb.CanonicalHostedZoneID

DNSName: !GetAtt IntApiElb.DNSName

- Name:

```
!Join [
  ".",
  ["api-int", !Ref ClusterName, !Join ["", [!Ref HostedZoneName, "."]]],
]
```

Type: A

AliasTarget:

HostedZoneId: !GetAtt IntApiElb.CanonicalHostedZoneID

DNSName: !GetAtt IntApiElb.DNSName

ExternalApiListener:

Type: AWS::ElasticLoadBalancingV2::Listener

Properties:

DefaultActions:

- Type: forward

TargetGroupArn:

Ref: ExternalApiTargetGroup

LoadBalancerArn:

Ref: ExtApiElb

Port: 6443

Protocol: TCP

ExternalApiTargetGroup:

Type: AWS::ElasticLoadBalancingV2::TargetGroup

Properties:

HealthCheckIntervalSeconds: 10

HealthCheckPath: "/readyz"

HealthCheckPort: 6443

HealthCheckProtocol: HTTPS

HealthyThresholdCount: 2

UnhealthyThresholdCount: 2

Port: 6443

Protocol: TCP

TargetType: ip

VpcId:

Ref: VpcId

TargetGroupAttributes:

- Key: deregistration_delay.timeout_seconds

Value: 60

InternalApiListener:

Type: AWS::ElasticLoadBalancingV2::Listener

Properties:

DefaultActions:
 - Type: forward
 TargetGroupArn:
 Ref: InternalApiTargetGroup
 LoadBalancerArn:
 Ref: IntApiElb
 Port: 6443
 Protocol: TCP

InternalApiTargetGroup:
 Type: AWS::ElasticLoadBalancingV2::TargetGroup
 Properties:
 HealthCheckIntervalSeconds: 10
 HealthCheckPath: "/readyz"
 HealthCheckPort: 6443
 HealthCheckProtocol: HTTPS
 HealthyThresholdCount: 2
 UnhealthyThresholdCount: 2
 Port: 6443
 Protocol: TCP
 TargetType: ip
 VpcId:
 Ref: VpcId
 TargetGroupAttributes:
 - Key: deregistration_delay.timeout_seconds
 Value: 60

InternalServiceInternalListener:
 Type: AWS::ElasticLoadBalancingV2::Listener
 Properties:
 DefaultActions:
 - Type: forward
 TargetGroupArn:
 Ref: InternalServiceTargetGroup
 LoadBalancerArn:
 Ref: IntApiElb
 Port: 22623
 Protocol: TCP

InternalServiceTargetGroup:
 Type: AWS::ElasticLoadBalancingV2::TargetGroup
 Properties:
 HealthCheckIntervalSeconds: 10
 HealthCheckPath: "/healthz"
 HealthCheckPort: 22623
 HealthCheckProtocol: HTTPS
 HealthyThresholdCount: 2
 UnhealthyThresholdCount: 2
 Port: 22623
 Protocol: TCP
 TargetType: ip
 VpcId:
 Ref: VpcId
 TargetGroupAttributes:
 - Key: deregistration_delay.timeout_seconds
 Value: 60

```

RegisterTargetLambdalamRole:
  Type: AWS::IAM::Role
  Properties:
    RoleName: !Join ["-", [!Ref InfrastructureName, "nlb", "lambda", "role"]]
    AssumeRolePolicyDocument:
      Version: "2012-10-17"
      Statement:
        - Effect: "Allow"
          Principal:
            Service:
              - "lambda.amazonaws.com"
          Action:
            - "sts:AssumeRole"
    Path: "/"
  Policies:
    - PolicyName: !Join ["-", [!Ref InfrastructureName, "master", "policy"]]
      PolicyDocument:
        Version: "2012-10-17"
        Statement:
          - Effect: "Allow"
            Action:
              [
                "elasticloadbalancing:RegisterTargets",
                "elasticloadbalancing:DeregisterTargets",
              ]
            Resource: !Ref InternalApiTargetGroup
          - Effect: "Allow"
            Action:
              [
                "elasticloadbalancing:RegisterTargets",
                "elasticloadbalancing:DeregisterTargets",
              ]
            Resource: !Ref InternalServiceTargetGroup
          - Effect: "Allow"
            Action:
              [
                "elasticloadbalancing:RegisterTargets",
                "elasticloadbalancing:DeregisterTargets",
              ]
            Resource: !Ref ExternalApiTargetGroup

```

```

RegisterNlbTargets:
  Type: "AWS::Lambda::Function"
  Properties:
    Handler: "index.handler"
    Role:
      Fn::GetAtt:
        - "RegisterTargetLambdalamRole"
        - "Arn"
  Code:
    ZipFile: |
      import json
      import boto3
      import cfnresponse
      def handler(event, context):

```

```

elb = boto3.client('elbv2')
if event['RequestType'] == 'Delete':
    elb.deregister_targets(TargetGroupArn=event['ResourceProperties']
['TargetArn'],Targets=[{'Id': event['ResourceProperties']['TargetIp']})
    elif event['RequestType'] == 'Create':
        elb.register_targets(TargetGroupArn=event['ResourceProperties']['TargetArn'],Targets=
[{'Id': event['ResourceProperties']['TargetIp']})
        responseData = {}
        cfnresponse.send(event, context, cfnresponse.SUCCESS, responseData,
event['ResourceProperties']['TargetArn']+event['ResourceProperties']['TargetIp'])
    Runtime: "python3.8"
    Timeout: 120

```

RegisterSubnetTagsLambdalamRole:

Type: AWS::IAM::Role

Properties:

RoleName: !Join ["-", [!Ref InfrastructureName, "subnet-tags-lambda-role"]]

AssumeRolePolicyDocument:

Version: "2012-10-17"

Statement:

- Effect: "Allow"

Principal:

Service:

- "lambda.amazonaws.com"

Action:

- "sts:AssumeRole"

Path: "/"

Policies:

- PolicyName: !Join ["-", [!Ref InfrastructureName, "subnet-tagging-policy"]]

PolicyDocument:

Version: "2012-10-17"

Statement:

- Effect: "Allow"

Action:

```

[
  "ec2:DeleteTags",
  "ec2:CreateTags"
]

```

Resource: "arn:aws:ec2:*:*:subnet/*"

- Effect: "Allow"

Action:

```

[
  "ec2:DescribeSubnets",
  "ec2:DescribeTags"
]

```

Resource: ""

RegisterSubnetTags:

Type: "AWS::Lambda::Function"

Properties:

Handler: "index.handler"

Role:

Fn::GetAtt:

- "RegisterSubnetTagsLambdalamRole"

- "Arn"

Code:

```

ZipFile: |
import json
import boto3
import cfnresponse
def handler(event, context):
    ec2_client = boto3.client('ec2')
    if event['RequestType'] == 'Delete':
        for subnet_id in event['ResourceProperties']['Subnets']:
            ec2_client.delete_tags(Resources=[subnet_id], Tags=[{'Key': 'kubernetes.io/cluster/' +
event['ResourceProperties']['InfrastructureName']}]);
    elif event['RequestType'] == 'Create':
        for subnet_id in event['ResourceProperties']['Subnets']:
            ec2_client.create_tags(Resources=[subnet_id], Tags=[{'Key': 'kubernetes.io/cluster/' +
event['ResourceProperties']['InfrastructureName'], 'Value': 'shared'}]);
        responseData = {}
        cfnresponse.send(event, context, cfnresponse.SUCCESS, responseData,
event['ResourceProperties']['InfrastructureName']+event['ResourceProperties']['Subnets'][0])
Runtime: "python3.8"
Timeout: 120

```

```

RegisterPublicSubnetTags:
Type: Custom::SubnetRegister
Properties:
ServiceToken: !GetAtt RegisterSubnetTags.Arn
InfrastructureName: !Ref InfrastructureName
Subnets: !Ref PublicSubnets

```

```

RegisterPrivateSubnetTags:
Type: Custom::SubnetRegister
Properties:
ServiceToken: !GetAtt RegisterSubnetTags.Arn
InfrastructureName: !Ref InfrastructureName
Subnets: !Ref PrivateSubnets

```

Outputs:

```

PrivateHostedZoneId:
Description: Hosted zone ID for the private DNS, which is required for private records.
Value: !Ref IntDns
ExternalApiLoadBalancerName:
Description: Full name of the external API load balancer.
Value: !GetAtt ExtApiElb.LoadBalancerFullName
InternalApiLoadBalancerName:
Description: Full name of the internal API load balancer.
Value: !GetAtt IntApiElb.LoadBalancerFullName
ApiServerDnsName:
Description: Full hostname of the API server, which is required for the Ignition config files.
Value: !Join [".", ["api-int", !Ref ClusterName, !Ref HostedZoneName]]
RegisterNlbIpTargetsLambda:
Description: Lambda ARN useful to help register or deregister IP targets for these load
balancers.
Value: !GetAtt RegisterNlbIpTargets.Arn
ExternalApiTargetGroupArn:
Description: ARN of the external API target group.
Value: !Ref ExternalApiTargetGroup
InternalApiTargetGroupArn:
Description: ARN of the internal API target group.

```


Value: !Ref InternalApiTargetGroup
 InternalServiceTargetGroupArn:
 Description: ARN of the internal service target group.
 Value: !Ref InternalServiceTargetGroup

IMPORTANT

If you are deploying your cluster to an AWS government or secret region, you must update the **InternalApiServerRecord** to use **CNAME** records. Records of type **ALIAS** are not supported for AWS government regions. For example:

Type: CNAME
 TTL: 10
 ResourceRecords:
 - !GetAtt IntApiElb.DNSName

Additional resources

- See [Listing public hosted zones](#) in the AWS documentation for more information about listing public hosted zones.

6.4.4.7. Creating security group and roles in AWS

You must create security groups and roles in Amazon Web Services (AWS) for your OpenShift Container Platform cluster to use.

You can use the provided CloudFormation template and a custom parameter file to create a stack of AWS resources. The stack represents the security groups and roles that your OpenShift Container Platform cluster requires.

NOTE

If you do not use the provided CloudFormation template to create your AWS infrastructure, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- You configured an AWS account.
- You added your AWS keys and region to your local AWS profile by running **aws configure**.
- You generated the Ignition config files for your cluster.
- You created and configured a VPC and associated subnets in AWS.

Procedure

1. Create a JSON file that contains the parameter values that the template requires:


[

```

{
  "ParameterKey": "InfrastructureName", 1
  "ParameterValue": "mycluster-<random_string>" 2
},
{
  "ParameterKey": "VpcCidr", 3
  "ParameterValue": "10.0.0.0/16" 4
},
{
  "ParameterKey": "PrivateSubnets", 5
  "ParameterValue": "subnet-<random_string>" 6
},
{
  "ParameterKey": "VpcId", 7
  "ParameterValue": "vpc-<random_string>" 8
}
]

```

- 1 The name for your cluster infrastructure that is encoded in your Ignition config files for the cluster.
 - 2 Specify the infrastructure name that you extracted from the Ignition config file metadata, which has the format **<cluster-name>-<random-string>**.
 - 3 The CIDR block for the VPC.
 - 4 Specify the CIDR block parameter that you used for the VPC that you defined in the form **x.x.x.x/16-24**.
 - 5 The private subnets that you created for your VPC.
 - 6 Specify the **PrivateSubnetIds** value from the output of the CloudFormation template for the VPC.
 - 7 The VPC that you created for the cluster.
 - 8 Specify the **VpcId** value from the output of the CloudFormation template for the VPC.
2. Copy the template from the **CloudFormation template for security objects** section of this topic and save it as a YAML file on your computer. This template describes the security groups and roles that your cluster requires.
 3. Launch the CloudFormation template to create a stack of AWS resources that represent the security groups and roles:



IMPORTANT

You must enter the command on a single line.

```

$ aws cloudformation create-stack --stack-name <name> 1
  --template-body file://<template>.yaml 2
  --parameters file://<parameters>.json 3
  --capabilities CAPABILITY_NAMED_IAM 4

```

- 1 **<name>** is the name for the CloudFormation stack, such as **cluster-sec**. You need the name of this stack if you remove the cluster.
- 2 **<template>** is the relative path to and name of the CloudFormation template YAML file that you saved.
- 3 **<parameters>** is the relative path to and name of the CloudFormation parameters JSON file.
- 4 You must explicitly declare the **CAPABILITY_NAMED_IAM** capability because the provided template creates some **AWS::IAM::Role** and **AWS::IAM::InstanceProfile** resources.

Example output

```
arn:aws:cloudformation:us-east-1:269333783861:stack/cluster-sec/03bd4210-2ed7-11eb-6d7a-13fc0b61e9db
```

4. Confirm that the template components exist:

```
$ aws cloudformation describe-stacks --stack-name <name>
```

After the **StackStatus** displays **CREATE_COMPLETE**, the output displays values for the following parameters. You must provide these parameter values to the other CloudFormation templates that you run to create your cluster:

MasterSecurityGroupID	Master Security Group ID
WorkerSecurityGroupID	Worker Security Group ID
MasterInstanceProfile	Master IAM Instance Profile
WorkerInstanceProfile	Worker IAM Instance Profile

6.4.4.7.1. CloudFormation template for security objects

You can use the following CloudFormation template to deploy the security objects that you need for your OpenShift Container Platform cluster.

Example 6.80. CloudFormation template for security objects

```
AWSTemplateFormatVersion: 2010-09-09
Description: Template for OpenShift Cluster Security Elements (Security Groups & IAM)
```

Parameters:

InfrastructureName:

AllowedPattern: `^[a-zA-Z][a-zA-Z0-9\-\]{0,26}$`

MaxLength: `27`

MinLength: `1`

ConstraintDescription: Infrastructure name must be alphanumeric, start with a letter, and have a maximum of `27` characters.

Description: A short, unique cluster ID used to tag cloud resources and identify items owned or used by the cluster.

Type: String

VpcCidr:

AllowedPattern: `^(((0-9){1,3}|1[0-9]{2}|2[0-4][0-9]|25[0-5])\.){3}((0-9){1,3}|1[0-9]{2}|2[0-4][0-9]|25[0-5])\.(1[6-9]|2[0-4])$`

ConstraintDescription: CIDR block parameter must be in the form `x.x.x.x/16-24`.

Default: `10.0.0.0/16`

Description: CIDR block for VPC.

Type: String

VpcId:

Description: The VPC-scoped resources will belong to this VPC.

Type: `AWS::EC2::VPC::Id`

PrivateSubnets:

Description: The internal subnets.

Type: `List<AWS::EC2::Subnet::Id>`

Metadata:

AWS::CloudFormation::Interface:

ParameterGroups:

- Label:

default: `"Cluster Information"`

Parameters:

- InfrastructureName

- Label:

default: `"Network Configuration"`

Parameters:

- VpcId

- VpcCidr

- PrivateSubnets

ParameterLabels:

InfrastructureName:

default: `"Infrastructure Name"`

VpcId:

default: `"VPC ID"`

VpcCidr:

default: `"VPC CIDR"`

PrivateSubnets:

default: `"Private Subnets"`

Resources:

MasterSecurityGroup:

Type: `AWS::EC2::SecurityGroup`

Properties:

GroupDescription: Cluster Master Security Group

SecurityGroupIngress:

- IpProtocol: `icmp`

FromPort: `0`

ToPort: 0
 CidrIp: !Ref VpcCidr
 - IpProtocol: tcp
 FromPort: 22
 ToPort: 22
 CidrIp: !Ref VpcCidr
 - IpProtocol: tcp
 ToPort: 6443
 FromPort: 6443
 CidrIp: !Ref VpcCidr
 - IpProtocol: tcp
 FromPort: 22623
 ToPort: 22623
 CidrIp: !Ref VpcCidr
 VpcId: !Ref VpcId

WorkerSecurityGroup:

Type: AWS::EC2::SecurityGroup
 Properties:
 GroupDescription: Cluster Worker Security Group
 SecurityGroupIngress:
 - IpProtocol: icmp
 FromPort: 0
 ToPort: 0
 CidrIp: !Ref VpcCidr
 - IpProtocol: tcp
 FromPort: 22
 ToPort: 22
 CidrIp: !Ref VpcCidr
 VpcId: !Ref VpcId

MasterIngressEtcd:

Type: AWS::EC2::SecurityGroupIngress
 Properties:
 GroupId: !GetAtt MasterSecurityGroup.GroupId
 SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId
 Description: etcd
 FromPort: 2379
 ToPort: 2380
 IpProtocol: tcp

MasterIngressVxlan:

Type: AWS::EC2::SecurityGroupIngress
 Properties:
 GroupId: !GetAtt MasterSecurityGroup.GroupId
 SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId
 Description: Vxlan packets
 FromPort: 4789
 ToPort: 4789
 IpProtocol: udp

MasterIngressWorkerVxlan:

Type: AWS::EC2::SecurityGroupIngress
 Properties:
 GroupId: !GetAtt MasterSecurityGroup.GroupId
 SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId

Description: Vxlan packets
FromPort: 4789
ToPort: 4789
IpProtocol: udp

MasterIngressGeneve:

Type: AWS::EC2::SecurityGroupIngress
Properties:
GroupId: !GetAtt MasterSecurityGroup.GroupId
SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId
Description: Geneve packets
FromPort: 6081
ToPort: 6081
IpProtocol: udp

MasterIngressWorkerGeneve:

Type: AWS::EC2::SecurityGroupIngress
Properties:
GroupId: !GetAtt MasterSecurityGroup.GroupId
SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId
Description: Geneve packets
FromPort: 6081
ToPort: 6081
IpProtocol: udp

MasterIngressIpsecIke:

Type: AWS::EC2::SecurityGroupIngress
Properties:
GroupId: !GetAtt MasterSecurityGroup.GroupId
SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId
Description: IPsec IKE packets
FromPort: 500
ToPort: 500
IpProtocol: udp

MasterIngressIpsecNat:

Type: AWS::EC2::SecurityGroupIngress
Properties:
GroupId: !GetAtt MasterSecurityGroup.GroupId
SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId
Description: IPsec NAT-T packets
FromPort: 4500
ToPort: 4500
IpProtocol: udp

MasterIngressIpsecEsp:

Type: AWS::EC2::SecurityGroupIngress
Properties:
GroupId: !GetAtt MasterSecurityGroup.GroupId
SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId
Description: IPsec ESP packets
IpProtocol: 50

MasterIngressWorkerIpsecIke:

Type: AWS::EC2::SecurityGroupIngress
Properties:

GroupId: !GetAtt MasterSecurityGroup.GroupId
 SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId
 Description: IPsec IKE packets
 FromPort: 500
 ToPort: 500
 IpProtocol: udp

MasterIngressWorkerIpsecNat:

Type: AWS::EC2::SecurityGroupIngress
 Properties:
 GroupId: !GetAtt MasterSecurityGroup.GroupId
 SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId
 Description: IPsec NAT-T packets
 FromPort: 4500
 ToPort: 4500
 IpProtocol: udp

MasterIngressWorkerIpsecEsp:

Type: AWS::EC2::SecurityGroupIngress
 Properties:
 GroupId: !GetAtt MasterSecurityGroup.GroupId
 SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId
 Description: IPsec ESP packets
 IpProtocol: 50

MasterIngressInternal:

Type: AWS::EC2::SecurityGroupIngress
 Properties:
 GroupId: !GetAtt MasterSecurityGroup.GroupId
 SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId
 Description: Internal cluster communication
 FromPort: 9000
 ToPort: 9999
 IpProtocol: tcp

MasterIngressWorkerInternal:

Type: AWS::EC2::SecurityGroupIngress
 Properties:
 GroupId: !GetAtt MasterSecurityGroup.GroupId
 SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId
 Description: Internal cluster communication
 FromPort: 9000
 ToPort: 9999
 IpProtocol: tcp

MasterIngressInternalUDP:

Type: AWS::EC2::SecurityGroupIngress
 Properties:
 GroupId: !GetAtt MasterSecurityGroup.GroupId
 SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId
 Description: Internal cluster communication
 FromPort: 9000
 ToPort: 9999
 IpProtocol: udp

MasterIngressWorkerInternalUDP:

Type: AWS::EC2::SecurityGroupIngress

Properties:

GroupId: !GetAtt MasterSecurityGroup.GroupId

SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId

Description: Internal cluster communication

FromPort: 9000

ToPort: 9999

IpProtocol: udp

MasterIngressKube:

Type: AWS::EC2::SecurityGroupIngress

Properties:

GroupId: !GetAtt MasterSecurityGroup.GroupId

SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId

Description: Kubernetes kubelet, scheduler and controller manager

FromPort: 10250

ToPort: 10259

IpProtocol: tcp

MasterIngressWorkerKube:

Type: AWS::EC2::SecurityGroupIngress

Properties:

GroupId: !GetAtt MasterSecurityGroup.GroupId

SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId

Description: Kubernetes kubelet, scheduler and controller manager

FromPort: 10250

ToPort: 10259

IpProtocol: tcp

MasterIngressIngressServices:

Type: AWS::EC2::SecurityGroupIngress

Properties:

GroupId: !GetAtt MasterSecurityGroup.GroupId

SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId

Description: Kubernetes ingress services

FromPort: 30000

ToPort: 32767

IpProtocol: tcp

MasterIngressWorkerIngressServices:

Type: AWS::EC2::SecurityGroupIngress

Properties:

GroupId: !GetAtt MasterSecurityGroup.GroupId

SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId

Description: Kubernetes ingress services

FromPort: 30000

ToPort: 32767

IpProtocol: tcp

MasterIngressIngressServicesUDP:

Type: AWS::EC2::SecurityGroupIngress

Properties:

GroupId: !GetAtt MasterSecurityGroup.GroupId

SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId

Description: Kubernetes ingress services

FromPort: 30000

ToPort: 32767
IpProtocol: udp

MasterIngressWorkerIngressServicesUDP:

Type: AWS::EC2::SecurityGroupIngress
Properties:
GroupId: !GetAtt MasterSecurityGroup.GroupId
SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId
Description: Kubernetes ingress services
FromPort: 30000
ToPort: 32767
IpProtocol: udp

WorkerIngressVxlan:

Type: AWS::EC2::SecurityGroupIngress
Properties:
GroupId: !GetAtt WorkerSecurityGroup.GroupId
SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId
Description: Vxlan packets
FromPort: 4789
ToPort: 4789
IpProtocol: udp

WorkerIngressMasterVxlan:

Type: AWS::EC2::SecurityGroupIngress
Properties:
GroupId: !GetAtt WorkerSecurityGroup.GroupId
SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId
Description: Vxlan packets
FromPort: 4789
ToPort: 4789
IpProtocol: udp

WorkerIngressGeneve:

Type: AWS::EC2::SecurityGroupIngress
Properties:
GroupId: !GetAtt WorkerSecurityGroup.GroupId
SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId
Description: Geneve packets
FromPort: 6081
ToPort: 6081
IpProtocol: udp

WorkerIngressMasterGeneve:

Type: AWS::EC2::SecurityGroupIngress
Properties:
GroupId: !GetAtt WorkerSecurityGroup.GroupId
SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId
Description: Geneve packets
FromPort: 6081
ToPort: 6081
IpProtocol: udp

WorkerIngressIpsecIke:

Type: AWS::EC2::SecurityGroupIngress
Properties:

GroupId: !GetAtt WorkerSecurityGroup.GroupId
SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId
Description: IPsec IKE packets
FromPort: 500
ToPort: 500
IpProtocol: udp

WorkerIngressIpsecNat:

Type: AWS::EC2::SecurityGroupIngress
Properties:
GroupId: !GetAtt WorkerSecurityGroup.GroupId
SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId
Description: IPsec NAT-T packets
FromPort: 4500
ToPort: 4500
IpProtocol: udp

WorkerIngressIpsecEsp:

Type: AWS::EC2::SecurityGroupIngress
Properties:
GroupId: !GetAtt WorkerSecurityGroup.GroupId
SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId
Description: IPsec ESP packets
IpProtocol: 50

WorkerIngressMasterIpsecIke:

Type: AWS::EC2::SecurityGroupIngress
Properties:
GroupId: !GetAtt WorkerSecurityGroup.GroupId
SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId
Description: IPsec IKE packets
FromPort: 500
ToPort: 500
IpProtocol: udp

WorkerIngressMasterIpsecNat:

Type: AWS::EC2::SecurityGroupIngress
Properties:
GroupId: !GetAtt WorkerSecurityGroup.GroupId
SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId
Description: IPsec NAT-T packets
FromPort: 4500
ToPort: 4500
IpProtocol: udp

WorkerIngressMasterIpsecEsp:

Type: AWS::EC2::SecurityGroupIngress
Properties:
GroupId: !GetAtt WorkerSecurityGroup.GroupId
SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId
Description: IPsec ESP packets
IpProtocol: 50

WorkerIngressInternal:

Type: AWS::EC2::SecurityGroupIngress
Properties:

GroupId: !GetAtt WorkerSecurityGroup.GroupId
SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId
Description: Internal cluster communication
FromPort: 9000
ToPort: 9999
IpProtocol: tcp

WorkerIngressMasterInternal:

Type: AWS::EC2::SecurityGroupIngress
Properties:
GroupId: !GetAtt WorkerSecurityGroup.GroupId
SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId
Description: Internal cluster communication
FromPort: 9000
ToPort: 9999
IpProtocol: tcp

WorkerIngressInternalUDP:

Type: AWS::EC2::SecurityGroupIngress
Properties:
GroupId: !GetAtt WorkerSecurityGroup.GroupId
SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId
Description: Internal cluster communication
FromPort: 9000
ToPort: 9999
IpProtocol: udp

WorkerIngressMasterInternalUDP:

Type: AWS::EC2::SecurityGroupIngress
Properties:
GroupId: !GetAtt WorkerSecurityGroup.GroupId
SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId
Description: Internal cluster communication
FromPort: 9000
ToPort: 9999
IpProtocol: udp

WorkerIngressKube:

Type: AWS::EC2::SecurityGroupIngress
Properties:
GroupId: !GetAtt WorkerSecurityGroup.GroupId
SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId
Description: Kubernetes secure kubelet port
FromPort: 10250
ToPort: 10250
IpProtocol: tcp

WorkerIngressWorkerKube:

Type: AWS::EC2::SecurityGroupIngress
Properties:
GroupId: !GetAtt WorkerSecurityGroup.GroupId
SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId
Description: Internal Kubernetes communication
FromPort: 10250
ToPort: 10250
IpProtocol: tcp

WorkerIngressIngressServices:

Type: AWS::EC2::SecurityGroupIngress

Properties:

GroupId: !GetAtt WorkerSecurityGroup.GroupId

SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId

Description: Kubernetes ingress services

FromPort: 30000

ToPort: 32767

IpProtocol: tcp

WorkerIngressMasterIngressServices:

Type: AWS::EC2::SecurityGroupIngress

Properties:

GroupId: !GetAtt WorkerSecurityGroup.GroupId

SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId

Description: Kubernetes ingress services

FromPort: 30000

ToPort: 32767

IpProtocol: tcp

WorkerIngressIngressServicesUDP:

Type: AWS::EC2::SecurityGroupIngress

Properties:

GroupId: !GetAtt WorkerSecurityGroup.GroupId

SourceSecurityGroupId: !GetAtt WorkerSecurityGroup.GroupId

Description: Kubernetes ingress services

FromPort: 30000

ToPort: 32767

IpProtocol: udp

WorkerIngressMasterIngressServicesUDP:

Type: AWS::EC2::SecurityGroupIngress

Properties:

GroupId: !GetAtt WorkerSecurityGroup.GroupId

SourceSecurityGroupId: !GetAtt MasterSecurityGroup.GroupId

Description: Kubernetes ingress services

FromPort: 30000

ToPort: 32767

IpProtocol: udp

MasterIamRole:

Type: AWS::IAM::Role

Properties:

AssumeRolePolicyDocument:

Version: "2012-10-17"

Statement:

- Effect: "Allow"

Principal:

Service:

- "ec2.amazonaws.com"

Action:

- "sts:AssumeRole"

Policies:

- PolicyName: !Join ["-", [!Ref InfrastructureName, "master", "policy"]]

PolicyDocument:

Version: "2012-10-17"

Statement:

- Effect: "Allow"

Action:

- "ec2:AttachVolume"
- "ec2:AuthorizeSecurityGroupIngress"
- "ec2:CreateSecurityGroup"
- "ec2:CreateTags"
- "ec2:CreateVolume"
- "ec2>DeleteSecurityGroup"
- "ec2>DeleteVolume"
- "ec2:Describe*"
- "ec2:DetachVolume"
- "ec2:ModifyInstanceAttribute"
- "ec2:ModifyVolume"
- "ec2:RevokeSecurityGroupIngress"
- "elasticloadbalancing:AddTags"
- "elasticloadbalancing:AttachLoadBalancerToSubnets"
- "elasticloadbalancing:ApplySecurityGroupsToLoadBalancer"
- "elasticloadbalancing:CreateListener"
- "elasticloadbalancing:CreateLoadBalancer"
- "elasticloadbalancing:CreateLoadBalancerPolicy"
- "elasticloadbalancing:CreateLoadBalancerListeners"
- "elasticloadbalancing:CreateTargetGroup"
- "elasticloadbalancing:ConfigureHealthCheck"
- "elasticloadbalancing>DeleteListener"
- "elasticloadbalancing>DeleteLoadBalancer"
- "elasticloadbalancing>DeleteLoadBalancerListeners"
- "elasticloadbalancing>DeleteTargetGroup"
- "elasticloadbalancing:DeregisterInstancesFromLoadBalancer"
- "elasticloadbalancing:DeregisterTargets"
- "elasticloadbalancing:Describe*"
- "elasticloadbalancing:DetachLoadBalancerFromSubnets"
- "elasticloadbalancing:ModifyListener"
- "elasticloadbalancing:ModifyLoadBalancerAttributes"
- "elasticloadbalancing:ModifyTargetGroup"
- "elasticloadbalancing:ModifyTargetGroupAttributes"
- "elasticloadbalancing:RegisterInstancesWithLoadBalancer"
- "elasticloadbalancing:RegisterTargets"
- "elasticloadbalancing:SetLoadBalancerPoliciesForBackendServer"
- "elasticloadbalancing:SetLoadBalancerPoliciesOfListener"
- "kms:DescribeKey"

Resource: "*"

MasterInstanceProfile:

Type: "AWS::IAM::InstanceProfile"

Properties:

Roles:

- Ref: "MasterIamRole"

WorkerIamRole:

Type: AWS::IAM::Role

Properties:

AssumeRolePolicyDocument:

Version: "2012-10-17"

Statement:

```

- Effect: "Allow"
  Principal:
    Service:
      - "ec2.amazonaws.com"
  Action:
    - "sts:AssumeRole"
Policies:
- PolicyName: !Join ["-", [!Ref InfrastructureName, "worker", "policy"]]
  PolicyDocument:
    Version: "2012-10-17"
    Statement:
      - Effect: "Allow"
        Action:
          - "ec2:DescribeInstances"
          - "ec2:DescribeRegions"
        Resource: "*"

```

```

WorkerInstanceProfile:
  Type: "AWS::IAM::InstanceProfile"
  Properties:
    Roles:
      - Ref: "WorkerIamRole"

```

```

Outputs:
MasterSecurityGroupId:
  Description: Master Security Group ID
  Value: !GetAtt MasterSecurityGroup.GroupId

```

```

WorkerSecurityGroupId:
  Description: Worker Security Group ID
  Value: !GetAtt WorkerSecurityGroup.GroupId

```

```

MasterInstanceProfile:
  Description: Master IAM Instance Profile
  Value: !Ref MasterInstanceProfile

```

```

WorkerInstanceProfile:
  Description: Worker IAM Instance Profile
  Value: !Ref WorkerInstanceProfile

```

6.4.4.8. Accessing RHCOS AMIs with stream metadata

In OpenShift Container Platform, *stream metadata* provides standardized metadata about RHCOS in the JSON format and injects the metadata into the cluster. Stream metadata is a stable format that supports multiple architectures and is intended to be self-documenting for maintaining automation.

You can use the **coreos print-stream-json** sub-command of **openshift-install** to access information about the boot images in the stream metadata format. This command provides a method for printing stream metadata in a scriptable, machine-readable format.

For user-provisioned installations, the **openshift-install** binary contains references to the version of RHCOS boot images that are tested for use with OpenShift Container Platform, such as the AWS AMI.

Procedure

To parse the stream metadata, use one of the following methods:

- From a Go program, use the official **stream-metadata-go** library at <https://github.com/coreos/stream-metadata-go>. You can also view example code in the library.
- From another programming language, such as Python or Ruby, use the JSON library of your preferred programming language.
- From a command-line utility that handles JSON data, such as **jq**:
 - Print the current **x86_64** or **aarch64** AMI for an AWS region, such as **us-west-1**:

For x86_64

```
$ openshift-install coreos print-stream-json | jq -r
'.architectures.x86_64.images.aws.regions["us-west-1"].image'
```

Example output

```
ami-0d3e625f84626bbda
```

For aarch64

```
$ openshift-install coreos print-stream-json | jq -r
'.architectures.aarch64.images.aws.regions["us-west-1"].image'
```


Example output

```
ami-0af1d3b7fa5be2131
```

The output of this command is the AWS AMI ID for your designated architecture and the **us-west-1** region. The AMI must belong to the same region as the cluster.

6.4.4.9. RHCOS AMIs for the AWS infrastructure

Red Hat provides Red Hat Enterprise Linux CoreOS (RHCOS) AMIs that are valid for the various AWS regions and instance architectures that you can manually specify for your OpenShift Container Platform nodes.



NOTE

By importing your own AMI, you can also install to regions that do not have a published RHCOS AMI.

Table 6.25. x86_64 RHCOS AMIs

AWS zone	AWS AMI
af-south-1	ami-018de6b1be470b7e6
ap-east-1	ami-092f09a4c8aacf56a

AWS zone	AWS AMI
ap-northeast-1	ami-001c9fc85b77505f4
ap-northeast-2	ami-0a5d7f1966d88fba3
ap-northeast-3	ami-085a2726ceb4f5eeb
ap-south-1	ami-03f2c19071fb6eab3
ap-south-2	ami-0c82522890979d94b
ap-southeast-1	ami-06e5b9d16ead6c55c
ap-southeast-2	ami-0ccd429129fbf6bc0
ap-southeast-3	ami-096f9b4d41116d95c
ap-southeast-4	ami-0b511ab55f9f7d052
ca-central-1	ami-0e357287c651be1f2
ca-west-1	ami-0b1692e5776740901
eu-central-1	ami-08b13fb388ba5610d
eu-central-2	ami-04777298b55045ea9
eu-north-1	ami-05421993a4ee567b9
eu-south-1	ami-00959341869d34746
eu-south-2	ami-0a745b610522a87cc
eu-west-1	ami-0e48f85ec57bd7baa
eu-west-2	ami-0c015b1387acddc5e
eu-west-3	ami-01e466af77a95b93d
il-central-1	ami-0a1b706793b54d087
me-central-1	ami-09d58e8b2099ae5bc
me-south-1	ami-0f9c259bc4346599c

AWS zone	AWS AMI
sa-east-1	ami-06277517d966881a3
us-east-1	ami-05483066c3caaccf5
us-east-2	ami-0c704ce516493a479
us-gov-east-1	ami-0695b2cbdd1ec4d48
us-gov-west-1	ami-004404a0eaa427b39
us-west-1	ami-0aaa56637063be772
us-west-2	ami-0870c7a3d5ef4d2b3

Table 6.26. aarch64 RHCOS AMIs

AWS zone	AWS AMI
af-south-1	ami-0d96d447d3df14f4e
ap-east-1	ami-010236402dfba1717
ap-northeast-1	ami-08feeb6d9068bb12e
ap-northeast-2	ami-0772082c119147205
ap-northeast-3	ami-05bcbb13493d0516f
ap-south-1	ami-0034a539e8adcb817
ap-south-2	ami-0fc56b7c53c38ff70
ap-southeast-1	ami-0b50a0e92a58f3ad8
ap-southeast-2	ami-0697a9174ae206182
ap-southeast-3	ami-05ae309319a7ad504
ap-southeast-4	ami-00500414843baa657
ca-central-1	ami-05e76bf99d3b0d4c8
ca-west-1	ami-099fc953c221ec590

AWS zone	AWS AMI
eu-central-1	ami-0d2e2698884020b71
eu-central-2	ami-0a96ba21ddee01719
eu-north-1	ami-0ab8a1070d0b1d9a5
eu-south-1	ami-0d0465299a8b196c1
eu-south-2	ami-07d5aa3c6d468c738
eu-west-1	ami-08e7d209f26c09c80
eu-west-2	ami-0c8336d4e2c1f13cb
eu-west-3	ami-085d804b98acf0648
il-central-1	ami-02ce030e36aeb7643
me-central-1	ami-0dea3ac466040b77f
me-south-1	ami-02ef2a46887b9786d
sa-east-1	ami-0d19817d31b2399f9
us-east-1	ami-04859c888566a2c2f
us-east-2	ami-02f7e4a5dc66361bb
us-gov-east-1	ami-067ef782980ea96ad
us-gov-west-1	ami-0ce67540c1bb2319d
us-west-1	ami-0a09c5d2f287718a3
us-west-2	ami-06b94e64e595f6cee

6.4.4.10. Creating the bootstrap node in AWS

You must create the bootstrap node in Amazon Web Services (AWS) to use during OpenShift Container Platform cluster initialization. You do this by:

- Providing a location to serve the **bootstrap.ign** Ignition config file to your cluster. This file is located in your installation directory. The provided CloudFormation Template assumes that the Ignition config files for your cluster are served from an S3 bucket. If you choose to serve the files from another location, you must modify the templates.

- Using the provided CloudFormation template and a custom parameter file to create a stack of AWS resources. The stack represents the bootstrap node that your OpenShift Container Platform installation requires.



NOTE

If you do not use the provided CloudFormation template to create your bootstrap node, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- You configured an AWS account.
- You added your AWS keys and region to your local AWS profile by running **aws configure**.
- You generated the Ignition config files for your cluster.
- You created and configured a VPC and associated subnets in AWS.
- You created and configured DNS, load balancers, and listeners in AWS.
- You created the security groups and roles required for your cluster in AWS.

Procedure

1. Create the bucket by running the following command:

```
$ aws s3 mb s3://<cluster-name>-infra 1
```

- 1** **<cluster-name>-infra** is the bucket name. When creating the **install-config.yaml** file, replace **<cluster-name>** with the name specified for the cluster.

You must use a presigned URL for your S3 bucket, instead of the **s3://** schema, if you are:

- Deploying to a region that has endpoints that differ from the AWS SDK.
 - Deploying a proxy.
 - Providing your own custom endpoints.
2. Upload the **bootstrap.ign** Ignition config file to the bucket by running the following command:

```
$ aws s3 cp <installation_directory>/bootstrap.ign s3://<cluster-name>-infra/bootstrap.ign 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

3. Verify that the file uploaded by running the following command:

```
$ aws s3 ls s3://<cluster-name>-infra/
```

Example output

2019-04-03 16:15:16 314878 bootstrap.ign



NOTE

The bootstrap Ignition config file does contain secrets, like X.509 keys. The following steps provide basic security for the S3 bucket. To provide additional security, you can enable an S3 bucket policy to allow only certain users, such as the OpenShift IAM user, to access objects that the bucket contains. You can avoid S3 entirely and serve your bootstrap Ignition config file from any address that the bootstrap machine can reach.

4. Create a JSON file that contains the parameter values that the template requires:

```
[
  {
    "ParameterKey": "InfrastructureName", 1
    "ParameterValue": "mycluster-<random_string>" 2
  },
  {
    "ParameterKey": "RhcoshAmi", 3
    "ParameterValue": "ami-<random_string>" 4
  },
  {
    "ParameterKey": "AllowedBootstrapSshCidr", 5
    "ParameterValue": "0.0.0.0/0" 6
  },
  {
    "ParameterKey": "PublicSubnet", 7
    "ParameterValue": "subnet-<random_string>" 8
  },
  {
    "ParameterKey": "MasterSecurityGroupID", 9
    "ParameterValue": "sg-<random_string>" 10
  },
  {
    "ParameterKey": "VpcID", 11
    "ParameterValue": "vpc-<random_string>" 12
  },
  {
    "ParameterKey": "BootstrapIgnitionLocation", 13
    "ParameterValue": "s3://<bucket_name>/bootstrap.ign" 14
  },
  {
    "ParameterKey": "AutoRegisterELB", 15
    "ParameterValue": "yes" 16
  },
  {
    "ParameterKey": "RegisterNlbTargetsLambdaArn", 17
    "ParameterValue": "arn:aws:lambda:<aws_region>:<account_number>:function:
<dns_stack_name>-RegisterNlbTargets-<random_string>" 18
  }
]
```


```

    },
    {
      "ParameterKey": "ExternalApiTargetGroupArn", 19
      "ParameterValue": "arn:aws:elasticloadbalancing:<aws_region>:
<account_number>:targetgroup/<dns_stack_name>-Exter-<random_string>" 20
    },
    {
      "ParameterKey": "InternalApiTargetGroupArn", 21
      "ParameterValue": "arn:aws:elasticloadbalancing:<aws_region>:
<account_number>:targetgroup/<dns_stack_name>-Inter-<random_string>" 22
    },
    {
      "ParameterKey": "InternalServiceTargetGroupArn", 23
      "ParameterValue": "arn:aws:elasticloadbalancing:<aws_region>:
<account_number>:targetgroup/<dns_stack_name>-Inter-<random_string>" 24
    }
  ]

```

- 1** The name for your cluster infrastructure that is encoded in your Ignition config files for the cluster.
- 2** Specify the infrastructure name that you extracted from the Ignition config file metadata, which has the format **<cluster-name>-<random-string>**.
- 3** Current Red Hat Enterprise Linux CoreOS (RHCOS) AMI to use for the bootstrap node based on your selected architecture.
- 4** Specify a valid **AWS::EC2::Image::Id** value.
- 5** CIDR block to allow SSH access to the bootstrap node.
- 6** Specify a CIDR block in the format **x.x.x.x/16-24**.
- 7** The public subnet that is associated with your VPC to launch the bootstrap node into.
- 8** Specify the **PublicSubnetIds** value from the output of the CloudFormation template for the VPC.
- 9** The master security group ID (for registering temporary rules)
- 10** Specify the **MasterSecurityGroupID** value from the output of the CloudFormation template for the security group and roles.
- 11** The VPC created resources will belong to.
- 12** Specify the **VpcId** value from the output of the CloudFormation template for the VPC.
- 13** Location to fetch bootstrap Ignition config file from.
- 14** Specify the S3 bucket and file name in the form **s3://<bucket_name>/bootstrap.ign**.
- 15** Whether or not to register a network load balancer (NLB).
- 16** Specify **yes** or **no**. If you specify **yes**, you must provide a Lambda Amazon Resource Name (ARN) value.

17. The ARN for NLB IP target registration lambda group.
 18. Specify the **RegisterNlbTargetsLambda** value from the output of the CloudFormation template for DNS and load balancing. Use **arn:aws-us-gov** if deploying the cluster to an AWS GovCloud region.
 19. The ARN for external API load balancer target group.
 20. Specify the **ExternalApiTargetGroupArn** value from the output of the CloudFormation template for DNS and load balancing. Use **arn:aws-us-gov** if deploying the cluster to an AWS GovCloud region.
 21. The ARN for internal API load balancer target group.
 22. Specify the **InternalApiTargetGroupArn** value from the output of the CloudFormation template for DNS and load balancing. Use **arn:aws-us-gov** if deploying the cluster to an AWS GovCloud region.
 23. The ARN for internal service load balancer target group.
 24. Specify the **InternalServiceTargetGroupArn** value from the output of the CloudFormation template for DNS and load balancing. Use **arn:aws-us-gov** if deploying the cluster to an AWS GovCloud region.
5. Copy the template from the **CloudFormation template for the bootstrap machines** section of this topic and save it as a YAML file on your computer. This template describes the bootstrap machine that your cluster requires.
 6. Optional: If you are deploying the cluster with a proxy, you must update the ignition in the template to add the **ignition.config.proxy** fields. Additionally, if you have added the Amazon EC2, Elastic Load Balancing, and S3 VPC endpoints to your VPC, you must add these endpoints to the **noProxy** field.
 7. Launch the CloudFormation template to create a stack of AWS resources that represent the bootstrap node:



IMPORTANT

You must enter the command on a single line.

```
$ aws cloudformation create-stack --stack-name <name> 1
  --template-body file://<template>.yaml 2
  --parameters file://<parameters>.json 3
  --capabilities CAPABILITY_NAMED_IAM 4
```

1. **<name>** is the name for the CloudFormation stack, such as **cluster-bootstrap**. You need the name of this stack if you remove the cluster.
2. **<template>** is the relative path to and name of the CloudFormation template YAML file that you saved.
3. **<parameters>** is the relative path to and name of the CloudFormation parameters JSON file.

- 4 You must explicitly declare the **CAPABILITY_NAMED_IAM** capability because the provided template creates some **AWS::IAM::Role** and **AWS::IAM::InstanceProfile**

Example output

```
arn:aws:cloudformation:us-east-1:269333783861:stack/cluster-bootstrap/12944486-2add-11eb-9dee-12dace8e3a83
```

8. Confirm that the template components exist:

```
$ aws cloudformation describe-stacks --stack-name <name>
```

After the **StackStatus** displays **CREATE_COMPLETE**, the output displays values for the following parameters. You must provide these parameter values to the other CloudFormation templates that you run to create your cluster:

Bootstrap InstanceId	The bootstrap Instance ID.
Bootstrap PublicIp	The bootstrap node public IP address.
Bootstrap PrivateIp	The bootstrap node private IP address.

6.4.4.10.1. CloudFormation template for the bootstrap machine

You can use the following CloudFormation template to deploy the bootstrap machine that you need for your OpenShift Container Platform cluster.

Example 6.81. CloudFormation template for the bootstrap machine

```
AWSTemplateFormatVersion: 2010-09-09
Description: Template for OpenShift Cluster Bootstrap (EC2 Instance, Security Groups and IAM)

Parameters:
  InfrastructureName:
    AllowedPattern: ^([a-zA-Z][a-zA-Z0-9\-\_]{0,26})$
    MaxLength: 27
    MinLength: 1
    ConstraintDescription: Infrastructure name must be alphanumeric, start with a letter, and have a maximum of 27 characters.
    Description: A short, unique cluster ID used to tag cloud resources and identify items owned or used by the cluster.
    Type: String
  RhcosAmi:
    Description: Current Red Hat Enterprise Linux CoreOS AMI to use for bootstrap.
    Type: AWS::EC2::Image::Id
  AllowedBootstrapSshCidr:
    AllowedPattern: ^((([0-9]{1,3}[0-9]{1,3}[0-9]{1,3}[0-9]{1,3})|([0-9]{1,3}[0-9]{1,3}[0-9]{1,3})|([0-9]{1,3}[0-9]{1,3})|([0-9]{1,3})|([0-9]{1,3}))\.(([0-9]{1,3}[0-9]{1,3}[0-9]{1,3}|[0-9]{1,3}))\.(([0-9]{1,3}[0-9]{1,3}[0-9]{1,3}|[0-9]{1,3}))\.(([0-9]{1,3}[0-9]{1,3}[0-9]{1,3}|[0-9]{1,3})))$
    ConstraintDescription: CIDR block parameter must be in the form x.x.x.x/0-32.
```

Default: `0.0.0.0/0`

Description: CIDR block to allow SSH access to the bootstrap node.

Type: String

PublicSubnet:

Description: The public subnet to launch the bootstrap node into.

Type: `AWS::EC2::Subnet::Id`

MasterSecurityGroupId:

Description: The master security group ID for registering temporary rules.

Type: `AWS::EC2::SecurityGroup::Id`

VpcId:

Description: The VPC-scoped resources will belong to this VPC.

Type: `AWS::EC2::VPC::Id`

BootstrapIgnitionLocation:

Default: `s3://my-s3-bucket/bootstrap.ign`

Description: Ignition config file location.

Type: String

AutoRegisterELB:

Default: `"yes"`

AllowedValues:

- `"yes"`

- `"no"`

Description: Do you want to invoke NLB registration, which requires a Lambda ARN parameter?

Type: String

RegisterNlbTargetsLambdaArn:

Description: ARN for NLB IP target registration lambda.

Type: String

ExternalApiTargetGroupArn:

Description: ARN for external API load balancer target group.

Type: String

InternalApiTargetGroupArn:

Description: ARN for internal API load balancer target group.

Type: String

InternalServiceTargetGroupArn:

Description: ARN for internal service load balancer target group.

Type: String

BootstrapInstanceType:

Description: Instance type for the bootstrap EC2 instance

Default: `"i3.large"`

Type: String

Metadata:

`AWS::CloudFormation::Interface`:

ParameterGroups:

- Label:

default: `"Cluster Information"`

Parameters:

- `InfrastructureName`

- Label:

default: `"Host Information"`

Parameters:

- `RhcosAmi`

- `BootstrapIgnitionLocation`

- `MasterSecurityGroupId`

- Label:

default: `"Network Configuration"`

Parameters:

- VpcId
- AllowedBootstrapSshCidr
- PublicSubnet
- Label:
 - default: "Load Balancer Automation"
- Parameters:
 - AutoRegisterELB
 - RegisterNlbTargetsLambdaArn
 - ExternalApiTargetGroupArn
 - InternalApiTargetGroupArn
 - InternalServiceTargetGroupArn
- ParameterLabels:
 - InfrastructureName:
 - default: "Infrastructure Name"
 - VpcId:
 - default: "VPC ID"
 - AllowedBootstrapSshCidr:
 - default: "Allowed SSH Source"
 - PublicSubnet:
 - default: "Public Subnet"
 - RhcosAmi:
 - default: "Red Hat Enterprise Linux CoreOS AMI ID"
 - BootstrapIgnitionLocation:
 - default: "Bootstrap Ignition Source"
 - MasterSecurityGroupId:
 - default: "Master Security Group ID"
 - AutoRegisterELB:
 - default: "Use Provided ELB Automation"

Conditions:

DoRegistration: !Equals ["yes", !Ref AutoRegisterELB]

Resources:

BootstrapIamRole:

Type: AWS::IAM::Role

Properties:

AssumeRolePolicyDocument:

Version: "2012-10-17"

Statement:

- Effect: "Allow"

Principal:

Service:

- "ec2.amazonaws.com"

Action:

- "sts:AssumeRole"

Path: "/"

Policies:

- PolicyName: !Join ["-", [!Ref InfrastructureName, "bootstrap", "policy"]]

PolicyDocument:

Version: "2012-10-17"

Statement:

- Effect: "Allow"

Action: "ec2:Describe*"

Resource: "*"

- Effect: "Allow"

Action: "ec2:AttachVolume"

```

    Resource: "*"
  - Effect: "Allow"
    Action: "ec2:DetachVolume"
    Resource: "*"
  - Effect: "Allow"
    Action: "s3:GetObject"
    Resource: "*"

```

BootstrapInstanceProfile:

```

Type: "AWS::IAM::InstanceProfile"
Properties:
  Path: "/"
  Roles:
  - Ref: "BootstrapIamRole"

```

BootstrapSecurityGroup:

```

Type: AWS::EC2::SecurityGroup
Properties:
  GroupDescription: Cluster Bootstrap Security Group
  SecurityGroupIngress:
  - IpProtocol: tcp
    FromPort: 22
    ToPort: 22
    CidrIp: !Ref AllowedBootstrapSshCidr
  - IpProtocol: tcp
    ToPort: 19531
    FromPort: 19531
    CidrIp: 0.0.0.0/0
  VpcId: !Ref VpcId

```

BootstrapInstance:

```

Type: AWS::EC2::Instance
Properties:
  ImageId: !Ref RHCOSAmi
  IamInstanceProfile: !Ref BootstrapInstanceProfile
  InstanceType: !Ref BootstrapInstanceType
  NetworkInterfaces:
  - AssociatePublicIpAddress: "true"
    DeviceIndex: "0"
    GroupSet:
    - !Ref "BootstrapSecurityGroup"
    - !Ref "MasterSecurityGroup"
    SubnetId: !Ref "PublicSubnet"
  UserData:
    Fn::Base64: !Sub
      - '{"ignition":{"config":{"replace":{"source":"${S3Loc}"},"version":"3.1.0"}}}'
      - {
          S3Loc: !Ref BootstrapIgnitionLocation
        }

```

RegisterBootstrapApiTarget:

```

Condition: DoRegistration
Type: Custom::NLBRegister
Properties:
  ServiceToken: !Ref RegisterNLBTargetsLambdaArn
  TargetArn: !Ref ExternalApiTargetGroupArn

```

```
TargetIp: !GetAtt BootstrapInstance.PrivateIp
```

```
RegisterBootstrapInternalApiTarget:
```

```
Condition: DoRegistration
```

```
Type: Custom::NLBRegister
```

```
Properties:
```

```
ServiceToken: !Ref RegisterNlbIpTargetsLambdaArn
```

```
TargetArn: !Ref InternalApiTargetGroupArn
```

```
TargetIp: !GetAtt BootstrapInstance.PrivateIp
```

```
RegisterBootstrapInternalServiceTarget:
```

```
Condition: DoRegistration
```

```
Type: Custom::NLBRegister
```

```
Properties:
```

```
ServiceToken: !Ref RegisterNlbIpTargetsLambdaArn
```

```
TargetArn: !Ref InternalServiceTargetGroupArn
```

```
TargetIp: !GetAtt BootstrapInstance.PrivateIp
```

```
Outputs:
```

```
BootstrapInstanceId:
```

```
Description: Bootstrap Instance ID.
```

```
Value: !Ref BootstrapInstance
```

```
BootstrapPublicIp:
```

```
Description: The bootstrap node public IP address.
```

```
Value: !GetAtt BootstrapInstance.PublicIp
```

```
BootstrapPrivateIp:
```

```
Description: The bootstrap node private IP address.
```

```
Value: !GetAtt BootstrapInstance.PrivateIp
```


Additional resources

- See [RHCOS AMIs for the AWS infrastructure](#) for details about the Red Hat Enterprise Linux CoreOS (RHCOS) AMIs for the AWS zones.

6.4.4.11. Creating the control plane machines in AWS

You must create the control plane machines in Amazon Web Services (AWS) that your cluster will use.

You can use the provided CloudFormation template and a custom parameter file to create a stack of AWS resources that represent the control plane nodes.



IMPORTANT

The CloudFormation template creates a stack that represents three control plane nodes.



NOTE

If you do not use the provided CloudFormation template to create your control plane nodes, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- You configured an AWS account.
- You added your AWS keys and region to your local AWS profile by running **aws configure**.
- You generated the Ignition config files for your cluster.
- You created and configured a VPC and associated subnets in AWS.
- You created and configured DNS, load balancers, and listeners in AWS.
- You created the security groups and roles required for your cluster in AWS.
- You created the bootstrap machine.

Procedure

1. Create a JSON file that contains the parameter values that the template requires:

```
[
  {
    "ParameterKey": "InfrastructureName", 1
    "ParameterValue": "mycluster-<random_string>" 2
  },
  {
    "ParameterKey": "RhcOsAmi", 3
    "ParameterValue": "ami-<random_string>" 4
  },
  {
    "ParameterKey": "AutoRegisterDNS", 5
    "ParameterValue": "yes" 6
  },
  {
    "ParameterKey": "PrivateHostedZoneId", 7
    "ParameterValue": "<random_string>" 8
  },
  {
    "ParameterKey": "PrivateHostedZoneName", 9
    "ParameterValue": "mycluster.example.com" 10
  },
  {
    "ParameterKey": "Master0Subnet", 11
    "ParameterValue": "subnet-<random_string>" 12
  },
  {
    "ParameterKey": "Master1Subnet", 13
    "ParameterValue": "subnet-<random_string>" 14
  },
  {
    "ParameterKey": "Master2Subnet", 15
    "ParameterValue": "subnet-<random_string>" 16
  },
  {
```

```

    "ParameterKey": "MasterSecurityGroup", 17
    "ParameterValue": "sg-<random_string>" 18
  },
  {
    "ParameterKey": "IgnitionLocation", 19
    "ParameterValue": "https://api-int.<cluster_name>.<domain_name>:22623/config/master"
  } 20
},
{
  "ParameterKey": "CertificateAuthorities", 21
  "ParameterValue": "data:text/plain;charset=utf-8;base64,ABC...xYz==" 22
},
{
  "ParameterKey": "MasterInstanceProfileName", 23
  "ParameterValue": "<roles_stack>-MasterInstanceProfile-<random_string>" 24
},
{
  "ParameterKey": "MasterInstanceType", 25
  "ParameterValue": "" 26
},
{
  "ParameterKey": "AutoRegisterELB", 27
  "ParameterValue": "yes" 28
},
{
  "ParameterKey": "RegisterNlbTargetsLambdaArn", 29
  "ParameterValue": "arn:aws:lambda:<aws_region>:<account_number>:function:
<dns_stack_name>-RegisterNlbTargets-<random_string>" 30
},
{
  "ParameterKey": "ExternalApiTargetGroupArn", 31
  "ParameterValue": "arn:aws:elasticloadbalancing:<aws_region>:
<account_number>:targetgroup/<dns_stack_name>-Exter-<random_string>" 32
},
{
  "ParameterKey": "InternalApiTargetGroupArn", 33
  "ParameterValue": "arn:aws:elasticloadbalancing:<aws_region>:
<account_number>:targetgroup/<dns_stack_name>-Inter-<random_string>" 34
},
{
  "ParameterKey": "InternalServiceTargetGroupArn", 35
  "ParameterValue": "arn:aws:elasticloadbalancing:<aws_region>:
<account_number>:targetgroup/<dns_stack_name>-Inter-<random_string>" 36
}
]

```

- 1 The name for your cluster infrastructure that is encoded in your Ignition config files for the cluster.
- 2 Specify the infrastructure name that you extracted from the Ignition config file metadata, which has the format **<cluster-name>-<random-string>**.
- 3 Current Red Hat Enterprise Linux CoreOS (RHCOS) AMI to use for the control plane

- 4 Specify an **AWS::EC2::Image::Id** value.
- 5 Whether or not to perform DNS etcd registration.
- 6 Specify **yes** or **no**. If you specify **yes**, you must provide hosted zone information.
- 7 The Route 53 private zone ID to register the etcd targets with.
- 8 Specify the **PrivateHostedZoneId** value from the output of the CloudFormation template for DNS and load balancing.
- 9 The Route 53 zone to register the targets with.
- 10 Specify **<cluster_name>.<domain_name>** where **<domain_name>** is the Route 53 base domain that you used when you generated **install-config.yaml** file for the cluster. Do not include the trailing period (.) that is displayed in the AWS console.
- 11 13 15 A subnet, preferably private, to launch the control plane machines on.
- 12 14 16 Specify a subnet from the **PrivateSubnets** value from the output of the CloudFormation template for DNS and load balancing.
- 17 The master security group ID to associate with control plane nodes.
- 18 Specify the **MasterSecurityGroupId** value from the output of the CloudFormation template for the security group and roles.
- 19 The location to fetch control plane Ignition config file from.
- 20 Specify the generated Ignition config file location, https://api-int.<cluster_name>.<domain_name>:22623/config/master.
- 21 The base64 encoded certificate authority string to use.
- 22 Specify the value from the **master.ign** file that is in the installation directory. This value is the long string with the format **data:text/plain;charset=utf-8;base64,ABC...xYz==**.
- 23 The IAM profile to associate with control plane nodes.
- 24 Specify the **MasterInstanceProfile** parameter value from the output of the CloudFormation template for the security group and roles.
- 25 The type of AWS instance to use for the control plane machines based on your selected architecture.
- 26 The instance type value corresponds to the minimum resource requirements for control plane machines. For example **m6i.xlarge** is a type for AMD64 and **m6g.xlarge** is a type for ARM64.
- 27 Whether or not to register a network load balancer (NLB).
- 28 Specify **yes** or **no**. If you specify **yes**, you must provide a Lambda Amazon Resource Name (ARN) value.
- 29 The ARN for NLB IP target registration lambda group.
- 30

Specify the **RegisterNlbTargetsLambda** value from the output of the CloudFormation template for DNS and load balancing. Use **arn:aws-us-gov** if deploying the cluster to an

- 31 The ARN for external API load balancer target group.
 - 32 Specify the **ExternalApiTargetGroupArn** value from the output of the CloudFormation template for DNS and load balancing. Use **arn:aws-us-gov** if deploying the cluster to an AWS GovCloud region.
 - 33 The ARN for internal API load balancer target group.
 - 34 Specify the **InternalApiTargetGroupArn** value from the output of the CloudFormation template for DNS and load balancing. Use **arn:aws-us-gov** if deploying the cluster to an AWS GovCloud region.
 - 35 The ARN for internal service load balancer target group.
 - 36 Specify the **InternalServiceTargetGroupArn** value from the output of the CloudFormation template for DNS and load balancing. Use **arn:aws-us-gov** if deploying the cluster to an AWS GovCloud region.
2. Copy the template from the **CloudFormation template for control plane machines** section of this topic and save it as a YAML file on your computer. This template describes the control plane machines that your cluster requires.
 3. If you specified an **m5** instance type as the value for **MasterInstanceType**, add that instance type to the **MasterInstanceType.AllowedValues** parameter in the CloudFormation template.
 4. Launch the CloudFormation template to create a stack of AWS resources that represent the control plane nodes:



IMPORTANT

You must enter the command on a single line.

```
$ aws cloudformation create-stack --stack-name <name> 1
  --template-body file://<template>.yaml 2
  --parameters file://<parameters>.json 3
```

- 1 **<name>** is the name for the CloudFormation stack, such as **cluster-control-plane**. You need the name of this stack if you remove the cluster.
- 2 **<template>** is the relative path to and name of the CloudFormation template YAML file that you saved.
- 3 **<parameters>** is the relative path to and name of the CloudFormation parameters JSON file.

Example output

```
arn:aws:cloudformation:us-east-1:269333783861:stack/cluster-control-plane/21c7e2b0-2ee2-11eb-c6f6-0aa34627df4b
```

**NOTE**

The CloudFormation template creates a stack that represents three control plane nodes.

- Confirm that the template components exist:

```
$ aws cloudformation describe-stacks --stack-name <name>
```

6.4.4.11.1. CloudFormation template for control plane machines

You can use the following CloudFormation template to deploy the control plane machines that you need for your OpenShift Container Platform cluster.

Example 6.82. CloudFormation template for control plane machines

```
AWSTemplateFormatVersion: 2010-09-09
Description: Template for OpenShift Cluster Node Launch (EC2 master instances)

Parameters:
  InfrastructureName:
    AllowedPattern: ^([a-zA-Z][a-zA-Z0-9\-\_]{0,26})$
    MaxLength: 27
    MinLength: 1
    ConstraintDescription: Infrastructure name must be alphanumeric, start with a letter, and have a
maximum of 27 characters.
    Description: A short, unique cluster ID used to tag nodes for the kubelet cloud provider.
    Type: String
  RhcosAmi:
    Description: Current Red Hat Enterprise Linux CoreOS AMI to use for bootstrap.
    Type: AWS::EC2::Image::Id
  AutoRegisterDNS:
    Default: ""
    Description: unused
    Type: String
  PrivateHostedZoneId:
    Default: ""
    Description: unused
    Type: String
  PrivateHostedZoneName:
    Default: ""
    Description: unused
    Type: String
  Master0Subnet:
    Description: The subnets, recommend private, to launch the master nodes into.
    Type: AWS::EC2::Subnet::Id
  Master1Subnet:
    Description: The subnets, recommend private, to launch the master nodes into.
    Type: AWS::EC2::Subnet::Id
  Master2Subnet:
    Description: The subnets, recommend private, to launch the master nodes into.
    Type: AWS::EC2::Subnet::Id
  MasterSecurityGroupId:
    Description: The master security group ID to associate with master nodes.
    Type: AWS::EC2::SecurityGroup::Id
```


IgnitionLocation:

Default: `https://api-int.$CLUSTER_NAME.$DOMAIN:22623/config/master`

Description: Ignition config file location.

Type: String

CertificateAuthorities:

Default: `data:text/plain;charset=utf-8;base64,ABC...xYz==`

Description: Base64 encoded certificate authority string to use.

Type: String

MasterInstanceProfileName:

Description: IAM profile to associate with master nodes.

Type: String

MasterInstanceType:

Default: `m5.xlarge`

Type: String

AutoRegisterELB:

Default: `"yes"`

AllowedValues:

- `"yes"`

- `"no"`

Description: Do you want to invoke NLB registration, which requires a Lambda ARN parameter?

Type: String

RegisterNLBIPTargetsLambdaArn:

Description: ARN for NLB IP target registration lambda. Supply the value from the cluster infrastructure or select `"no"` for AutoRegisterELB.

Type: String

ExternalApiTargetGroupArn:

Description: ARN for external API load balancer target group. Supply the value from the cluster infrastructure or select `"no"` for AutoRegisterELB.

Type: String

InternalApiTargetGroupArn:

Description: ARN for internal API load balancer target group. Supply the value from the cluster infrastructure or select `"no"` for AutoRegisterELB.

Type: String

InternalServiceTargetGroupArn:

Description: ARN for internal service load balancer target group. Supply the value from the cluster infrastructure or select `"no"` for AutoRegisterELB.

Type: String

Metadata:

AWS::CloudFormation::Interface:

ParameterGroups:

- **Label:**

default: `"Cluster Information"`

Parameters:

- **InfrastructureName**

- **Label:**

default: `"Host Information"`

Parameters:

- **MasterInstanceType**

- **RhcosAmi**

- **IgnitionLocation**

- **CertificateAuthorities**

- **MasterSecurityGroupID**

- **MasterInstanceProfileName**

- **Label:**

```
    default: "Network Configuration"
Parameters:
- VpcId
- AllowedBootstrapSshCidr
- Master0Subnet
- Master1Subnet
- Master2Subnet
- Label:
  default: "Load Balancer Automation"
Parameters:
- AutoRegisterELB
- RegisterNlbIpTargetsLambdaArn
- ExternalApiTargetGroupArn
- InternalApiTargetGroupArn
- InternalServiceTargetGroupArn
ParameterLabels:
InfrastructureName:
  default: "Infrastructure Name"
VpcId:
  default: "VPC ID"
Master0Subnet:
  default: "Master-0 Subnet"
Master1Subnet:
  default: "Master-1 Subnet"
Master2Subnet:
  default: "Master-2 Subnet"
MasterInstanceType:
  default: "Master Instance Type"
MasterInstanceProfileName:
  default: "Master Instance Profile Name"
RhcOsAmi:
  default: "Red Hat Enterprise Linux CoreOS AMI ID"
BootstrapIgnitionLocation:
  default: "Master Ignition Source"
CertificateAuthorities:
  default: "Ignition CA String"
MasterSecurityGroupId:
  default: "Master Security Group ID"
AutoRegisterELB:
  default: "Use Provided ELB Automation"

Conditions:
DoRegistration: !Equals ["yes", !Ref AutoRegisterELB]

Resources:
Master0:
  Type: AWS::EC2::Instance
  Properties:
    ImageId: !Ref RhcOsAmi
    BlockDeviceMappings:
      - DeviceName: /dev/xvda
    Ebs:
      VolumeSize: "120"
      VolumeType: "gp2"
    IamInstanceProfile: !Ref MasterInstanceProfileName
    InstanceType: !Ref MasterInstanceType
```

```

NetworkInterfaces:
- AssociatePublicIpAddress: "false"
  DeviceIndex: "0"
  GroupSet:
  - !Ref "MasterSecurityGroupId"
  SubnetId: !Ref "Master0Subnet"
UserData:
  Fn::Base64: !Sub
  - '{"ignition":{"config":{"merge":[{"source":"${SOURCE}"}]},"security":{"tls":
{"certificateAuthorities":[{"source":"${CA_BUNDLE}"}]}}, "version":"3.1.0"}}'
  - {
    SOURCE: !Ref IgnitionLocation,
    CA_BUNDLE: !Ref CertificateAuthorities,
  }
Tags:
- Key: !Join ["", ["kubernetes.io/cluster/", !Ref InfrastructureName]]
  Value: "shared"

```

```

RegisterMaster0:
  Condition: DoRegistration
  Type: Custom::NLBRegister
  Properties:
    ServiceToken: !Ref RegisterNlbTargetsLambdaArn
    TargetArn: !Ref ExternalApiTargetGroupArn
    TargetIp: !GetAtt Master0.PrivateIp

```

```

RegisterMaster0InternalApiTarget:
  Condition: DoRegistration
  Type: Custom::NLBRegister
  Properties:
    ServiceToken: !Ref RegisterNlbTargetsLambdaArn
    TargetArn: !Ref InternalApiTargetGroupArn
    TargetIp: !GetAtt Master0.PrivateIp

```

```

RegisterMaster0InternalServiceTarget:
  Condition: DoRegistration
  Type: Custom::NLBRegister
  Properties:
    ServiceToken: !Ref RegisterNlbTargetsLambdaArn
    TargetArn: !Ref InternalServiceTargetGroupArn
    TargetIp: !GetAtt Master0.PrivateIp

```

```

Master1:
  Type: AWS::EC2::Instance
  Properties:
    ImageId: !Ref RhcosAmi
    BlockDeviceMappings:
    - DeviceName: /dev/xvda
      Ebs:
        VolumeSize: "120"
        VolumeType: "gp2"
    IamInstanceProfile: !Ref MasterInstanceProfileName
    InstanceType: !Ref MasterInstanceType
    NetworkInterfaces:
    - AssociatePublicIpAddress: "false"
      DeviceIndex: "0"

```

```

GroupSet:
- !Ref "MasterSecurityGroupId"
SubnetId: !Ref "Master1Subnet"
UserData:
  Fn::Base64: !Sub
    - '{"ignition":{"config":{"merge":{"source":"${SOURCE}"},"security":{"tls":{"certificateAuthorities":{"source":"${CA_BUNDLE}"},"version":"3.1.0"}}}}'
    - {
      SOURCE: !Ref IgnitionLocation,
      CA_BUNDLE: !Ref CertificateAuthorities,
    }
  Tags:
    - Key: !Join [ "", ["kubernetes.io/cluster/", !Ref InfrastructureName] ]
      Value: "shared"

```

```

RegisterMaster1:
  Condition: DoRegistration
  Type: Custom::NLBRegister
  Properties:
    ServiceToken: !Ref RegisterNlbTargetsLambdaArn
    TargetArn: !Ref ExternalApiTargetGroupArn
    TargetIp: !GetAtt Master1.PrivateIp

```

```

RegisterMaster1InternalApiTarget:
  Condition: DoRegistration
  Type: Custom::NLBRegister
  Properties:
    ServiceToken: !Ref RegisterNlbTargetsLambdaArn
    TargetArn: !Ref InternalApiTargetGroupArn
    TargetIp: !GetAtt Master1.PrivateIp

```

```

RegisterMaster1InternalServiceTarget:
  Condition: DoRegistration
  Type: Custom::NLBRegister
  Properties:
    ServiceToken: !Ref RegisterNlbTargetsLambdaArn
    TargetArn: !Ref InternalServiceTargetGroupArn
    TargetIp: !GetAtt Master1.PrivateIp

```

```

Master2:
  Type: AWS::EC2::Instance
  Properties:
    ImageId: !Ref RhcosAmi
    BlockDeviceMappings:
      - DeviceName: /dev/xvda
        Ebs:
          VolumeSize: "120"
          VolumeType: "gp2"
    IamInstanceProfile: !Ref MasterInstanceProfileName
    InstanceType: !Ref MasterInstanceType
    NetworkInterfaces:
      - AssociatePublicIpAddress: "false"
        DeviceIndex: "0"
        GroupSet:
          - !Ref "MasterSecurityGroupId"
        SubnetId: !Ref "Master2Subnet"

```

```

UserData:
  Fn::Base64: !Sub
    - '{"ignition":{"config":{"merge":[{"source":"${SOURCE}"}]},"security":{"tls":
{"certificateAuthorities":[{"source":"${CA_BUNDLE}"}]},"version":"3.1.0"}}'
    - {
      SOURCE: !Ref IgnitionLocation,
      CA_BUNDLE: !Ref CertificateAuthorities,
    }
  Tags:
    - Key: !Join [ "", ["kubernetes.io/cluster/", !Ref InfrastructureName]]
      Value: "shared"

RegisterMaster2:
  Condition: DoRegistration
  Type: Custom::NLBRegister
  Properties:
    ServiceToken: !Ref RegisterNlbTargetsLambdaArn
    TargetArn: !Ref ExternalApiTargetGroupArn
    TargetIp: !GetAtt Master2.PrivateIp

RegisterMaster2InternalApiTarget:
  Condition: DoRegistration
  Type: Custom::NLBRegister
  Properties:
    ServiceToken: !Ref RegisterNlbTargetsLambdaArn
    TargetArn: !Ref InternalApiTargetGroupArn
    TargetIp: !GetAtt Master2.PrivateIp

RegisterMaster2InternalServiceTarget:
  Condition: DoRegistration
  Type: Custom::NLBRegister
  Properties:
    ServiceToken: !Ref RegisterNlbTargetsLambdaArn
    TargetArn: !Ref InternalServiceTargetGroupArn
    TargetIp: !GetAtt Master2.PrivateIp

Outputs:
  PrivateIPs:
    Description: The control-plane node private IP addresses.
    Value:
      !Join [
        "",
        [!GetAtt Master0.PrivateIp, !GetAtt Master1.PrivateIp, !GetAtt Master2.PrivateIp]
      ]

```

6.4.4.12. Creating the worker nodes in AWS

You can create worker nodes in Amazon Web Services (AWS) for your cluster to use.

You can use the provided CloudFormation template and a custom parameter file to create a stack of AWS resources that represent a worker node.



IMPORTANT

The CloudFormation template creates a stack that represents one worker node. You must create a stack for each worker node.



NOTE

If you do not use the provided CloudFormation template to create your worker nodes, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- You configured an AWS account.
- You added your AWS keys and region to your local AWS profile by running **aws configure**.
- You generated the Ignition config files for your cluster.
- You created and configured a VPC and associated subnets in AWS.
- You created and configured DNS, load balancers, and listeners in AWS.
- You created the security groups and roles required for your cluster in AWS.
- You created the bootstrap machine.
- You created the control plane machines.

Procedure

1. Create a JSON file that contains the parameter values that the CloudFormation template requires:

```
[
  {
    "ParameterKey": "InfrastructureName", 1
    "ParameterValue": "mycluster-<random_string>" 2
  },
  {
    "ParameterKey": "RhcocAmi", 3
    "ParameterValue": "ami-<random_string>" 4
  },
  {
    "ParameterKey": "Subnet", 5
    "ParameterValue": "subnet-<random_string>" 6
  },
  {
    "ParameterKey": "WorkerSecurityGroupID", 7
    "ParameterValue": "sg-<random_string>" 8
  },
  {
    "ParameterKey": "IgnitionLocation", 9
    "ParameterValue": "https://api-int.<cluster_name>.<domain_name>:22623/config/worker"
```

```


10
  },
  {
    "ParameterKey": "CertificateAuthorities", 11
    "ParameterValue": "" 12
  },
  {
    "ParameterKey": "WorkerInstanceProfileName", 13
    "ParameterValue": "" 14
  },
  {
    "ParameterKey": "WorkerInstanceType", 15
    "ParameterValue": "" 16
  }
}
]

```

- 1 The name for your cluster infrastructure that is encoded in your Ignition config files for the cluster.
- 2 Specify the infrastructure name that you extracted from the Ignition config file metadata, which has the format **<cluster-name>-<random-string>**.
- 3 Current Red Hat Enterprise Linux CoreOS (RHCOS) AMI to use for the worker nodes based on your selected architecture.
- 4 Specify an **AWS::EC2::Image::Id** value.
- 5 A subnet, preferably private, to start the worker nodes on.
- 6 Specify a subnet from the **PrivateSubnets** value from the output of the CloudFormation template for DNS and load balancing.
- 7 The worker security group ID to associate with worker nodes.
- 8 Specify the **WorkerSecurityGroupId** value from the output of the CloudFormation template for the security group and roles.
- 9 The location to fetch the bootstrap Ignition config file from.
- 10 Specify the generated Ignition config location, https://api-int.<cluster_name>.<domain_name>:22623/config/worker.
- 11 Base64 encoded certificate authority string to use.
- 12 Specify the value from the **worker.ign** file that is in the installation directory. This value is the long string with the format **data:text/plain;charset=utf-8;base64,ABC...xYz==**.
- 13 The IAM profile to associate with worker nodes.
- 14 Specify the **WorkerInstanceProfile** parameter value from the output of the CloudFormation template for the security group and roles.
- 15 The type of AWS instance to use for the compute machines based on your selected architecture.
- 16

The instance type value corresponds to the minimum resource requirements for compute machines. For example **m6i.large** is a type for AMD64 and **m6g.large** is a type for ARM64.

- Copy the template from the **CloudFormation template for worker machines** section of this topic and save it as a YAML file on your computer. This template describes the networking objects and load balancers that your cluster requires.
- Optional: If you specified an **m5** instance type as the value for **WorkerInstanceType**, add that instance type to the **WorkerInstanceType.AllowedValues** parameter in the CloudFormation template.
- Optional: If you are deploying with an AWS Marketplace image, update the **Worker0.type.properties.ImageID** parameter with the AMI ID that you obtained from your subscription.
- Use the CloudFormation template to create a stack of AWS resources that represent a worker node:



IMPORTANT

You must enter the command on a single line.

```
$ aws cloudformation create-stack --stack-name <name> 1
  --template-body file://<template>.yaml \ 2
  --parameters file://<parameters>.json 3
```

- <name>** is the name for the CloudFormation stack, such as **cluster-worker-1**. You need the name of this stack if you remove the cluster.
- <template>** is the relative path to and name of the CloudFormation template YAML file that you saved.
- <parameters>** is the relative path to and name of the CloudFormation parameters JSON file.

Example output

```
arn:aws:cloudformation:us-east-1:269333783861:stack/cluster-worker-1/729ee301-1c2a-11eb-348f-sd9888c65b59
```




NOTE

The CloudFormation template creates a stack that represents one worker node.

- Confirm that the template components exist:

```
$ aws cloudformation describe-stacks --stack-name <name>
```

- Continue to create worker stacks until you have created enough worker machines for your cluster. You can create additional worker stacks by referencing the same template and parameter files and specifying a different stack name.

**IMPORTANT**

You must create at least two worker machines, so you must create at least two stacks that use this CloudFormation template.

6.4.4.12.1. CloudFormation template for worker machines

You can use the following CloudFormation template to deploy the worker machines that you need for your OpenShift Container Platform cluster.

Example 6.83. CloudFormation template for worker machines

```

AWSTemplateFormatVersion: 2010-09-09
Description: Template for OpenShift Cluster Node Launch (EC2 worker instance)

Parameters:
  InfrastructureName:
    AllowedPattern: ^([a-zA-Z][a-zA-Z0-9\-\]{0,26})$
    MaxLength: 27
    MinLength: 1
    ConstraintDescription: Infrastructure name must be alphanumeric, start with a letter, and have a
    maximum of 27 characters.
    Description: A short, unique cluster ID used to tag nodes for the kubelet cloud provider.
    Type: String
  RhcosAmi:
    Description: Current Red Hat Enterprise Linux CoreOS AMI to use for bootstrap.
    Type: AWS::EC2::Image::Id
  Subnet:
    Description: The subnets, recommend private, to launch the worker nodes into.
    Type: AWS::EC2::Subnet::Id
  WorkerSecurityGroupId:
    Description: The worker security group ID to associate with worker nodes.
    Type: AWS::EC2::SecurityGroup::Id
  IgnitionLocation:
    Default: https://api-int.$CLUSTER_NAME.$DOMAIN:22623/config/worker
    Description: Ignition config file location.
    Type: String
  CertificateAuthorities:
    Default: data:text/plain;charset=utf-8;base64,ABC...xYz==
    Description: Base64 encoded certificate authority string to use.
    Type: String
  WorkerInstanceProfileName:
    Description: IAM profile to associate with worker nodes.
    Type: String
  WorkerInstanceType:
    Default: m5.large
    Type: String

Metadata:
  AWS::CloudFormation::Interface:
    ParameterGroups:
      - Label:
          default: "Cluster Information"
        Parameters:
          - InfrastructureName
          - Label:

```

```

    default: "Host Information"
Parameters:
- WorkerInstanceType
- RhcosAmi
- IgnitionLocation
- CertificateAuthorities
- WorkerSecurityGroupld
- WorkerInstanceProfileName
- Label:
  default: "Network Configuration"
Parameters:
- Subnet
ParameterLabels:
Subnet:
  default: "Subnet"
InfrastructureName:
  default: "Infrastructure Name"
WorkerInstanceType:
  default: "Worker Instance Type"
WorkerInstanceProfileName:
  default: "Worker Instance Profile Name"
RhcosAmi:
  default: "Red Hat Enterprise Linux CoreOS AMI ID"
IgnitionLocation:
  default: "Worker Ignition Source"
CertificateAuthorities:
  default: "Ignition CA String"
WorkerSecurityGroupld:
  default: "Worker Security Group ID"

Resources:
Worker0:
  Type: AWS::EC2::Instance
  Properties:
    ImageId: !Ref RhcosAmi
    BlockDeviceMappings:
      - DeviceName: /dev/xvda
        Ebs:
          VolumeSize: "120"
          VolumeType: "gp2"
    IamInstanceProfile: !Ref WorkerInstanceProfileName
    InstanceType: !Ref WorkerInstanceType
    NetworkInterfaces:
      - AssociatePublicIpAddress: "false"
        DeviceIndex: "0"
        GroupSet:
          - !Ref "WorkerSecurityGroupld"
        SubnetId: !Ref "Subnet"
    UserData:
      Fn::Base64: !Sub
        - '{"ignition":{"config":{"merge":[{"source":"${SOURCE}"]},"security":{"tls":
{"certificateAuthorities":[{"source":"${CA_BUNDLE}"]},"version":"3.1.0"}}}}'
        - {
          SOURCE: !Ref IgnitionLocation,
          CA_BUNDLE: !Ref CertificateAuthorities,
        }

```

```
Tags:
- Key: !Join [ "", ["kubernetes.io/cluster/", !Ref InfrastructureName]]
  Value: "shared"
```

```
Outputs:
PrivateIP:
  Description: The compute node private IP address.
  Value: !GetAtt Worker0.PrivateIp
```

6.4.4.13. Initializing the bootstrap sequence on AWS with user-provisioned infrastructure

After you create all of the required infrastructure in Amazon Web Services (AWS), you can start the bootstrap sequence that initializes the OpenShift Container Platform control plane.

Prerequisites

- You configured an AWS account.
- You added your AWS keys and region to your local AWS profile by running **aws configure**.
- You generated the Ignition config files for your cluster.
- You created and configured a VPC and associated subnets in AWS.
- You created and configured DNS, load balancers, and listeners in AWS.
- You created the security groups and roles required for your cluster in AWS.
- You created the bootstrap machine.
- You created the control plane machines.
- You created the worker nodes.

Procedure

1. Change to the directory that contains the installation program and start the bootstrap process that initializes the OpenShift Container Platform control plane:

```
$ ./openshift-install wait-for bootstrap-complete --dir <installation_directory> \ 1
--log-level=info 2
```

1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Example output

```
INFO Waiting up to 20m0s for the Kubernetes API at
https://api.mycluster.example.com:6443...
INFO API v1.29.4 up
```

```
INFO Waiting up to 30m0s for bootstrapping to complete...
INFO It is now safe to remove the bootstrap resources
INFO Time elapsed: 1s
```

If the command exits without a **FATAL** warning, your OpenShift Container Platform control plane has initialized.



NOTE

After the control plane initializes, it sets up the compute nodes and installs additional services in the form of Operators.

Additional resources

- See [Monitoring installation progress](#) for details about monitoring the installation, bootstrap, and control plane logs as an OpenShift Container Platform installation progresses.
- See [Gathering bootstrap node diagnostic data](#) for information about troubleshooting issues related to the bootstrap process.

6.4.4.14. Approving the certificate signing requests for your machines

When you add machines to a cluster, two pending certificate signing requests (CSRs) are generated for each machine that you added. You must confirm that these CSRs are approved or, if necessary, approve them yourself. The client requests must be approved first, followed by the server requests.

Prerequisites

- You added machines to your cluster.

Procedure

1. Confirm that the cluster recognizes the machines:

```
$ oc get nodes
```

Example output

```
NAME      STATUS  ROLES  AGE  VERSION
master-0  Ready   master 63m  v1.29.4
master-1  Ready   master 63m  v1.29.4
master-2  Ready   master 64m  v1.29.4
```

The output lists all of the machines that you created.



NOTE

The preceding output might not include the compute nodes, also known as worker nodes, until some CSRs are approved.

2. Review the pending CSRs and ensure that you see the client requests with the **Pending** or **Approved** status for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

```
NAME      AGE   REQUESTOR                                     CONDITION
csr-8b2br 15m   system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper Pending
csr-8vnps 15m   system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper Pending
...
```

In this example, two machines are joining the cluster. You might see more approved CSRs in the list.

- If the CSRs were not approved, after all of the pending CSRs for the machines you added are in **Pending** status, approve the CSRs for your cluster machines:



NOTE

Because the CSRs rotate automatically, approve your CSRs within an hour of adding the machines to the cluster. If you do not approve them within an hour, the certificates will rotate, and more than two certificates will be present for each node. You must approve all of these certificates. After the client CSR is approved, the Kubelet creates a secondary CSR for the serving certificate, which requires manual approval. Then, subsequent serving certificate renewal requests are automatically approved by the **machine-approver** if the Kubelet requests a new certificate with identical parameters.



NOTE

For clusters running on platforms that are not machine API enabled, such as bare metal and other user-provisioned infrastructure, you must implement a method of automatically approving the kubelet serving certificate requests (CSRs). If a request is not approved, then the **oc exec**, **oc rsh**, and **oc logs** commands cannot succeed, because a serving certificate is required when the API server connects to the kubelet. Any operation that contacts the Kubelet endpoint requires this certificate approval to be in place. The method must watch for new CSRs, confirm that the CSR was submitted by the **node-bootstrapper** service account in the **system:node** or **system:admin** groups, and confirm the identity of the node.


- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{\n"}\n{{end}}\n{{end}}' | xargs --no-run-if-empty oc adm certificate approve
```

**NOTE**

Some Operators might not become available until some CSRs are approved.

- Now that your client requests are approved, you must review the server requests for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

```
NAME      AGE   REQUESTOR                                     CONDITION
csr-bfd72 5m26s system:node:ip-10-0-50-126.us-east-2.compute.internal
Pending
csr-c57lv 5m26s system:node:ip-10-0-95-157.us-east-2.compute.internal
Pending
...
```

- If the remaining CSRs are not approved, and are in the **Pending** status, approve the CSRs for your cluster machines:

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

1 **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:


```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{"\n"}\n{{end}}{{end}}' | xargs oc adm certificate approve
```

- After all client and server CSRs have been approved, the machines have the **Ready** status. Verify this by running the following command:

```
$ oc get nodes
```

Example output

```
NAME      STATUS  ROLES  AGE  VERSION
master-0  Ready   master 73m  v1.29.4
master-1  Ready   master 73m  v1.29.4
master-2  Ready   master 74m  v1.29.4
worker-0  Ready   worker 11m  v1.29.4
worker-1  Ready   worker 11m  v1.29.4
```

**NOTE**

It can take a few minutes after approval of the server CSRs for the machines to transition to the **Ready** status.

Additional information

- For more information on CSRs, see [Certificate Signing Requests](#).

6.4.4.15. Initial Operator configuration

After the control plane initializes, you must immediately configure some Operators so that they all become available.

Prerequisites

- Your control plane has initialized.

Procedure

1. Watch the cluster components come online:

```
$ watch -n5 oc get clusteroperators
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.16.0	True	False	False	19m
baremetal	4.16.0	True	False	False	37m
cloud-credential	4.16.0	True	False	False	40m
cluster-autoscaler	4.16.0	True	False	False	37m
config-operator	4.16.0	True	False	False	38m
console	4.16.0	True	False	False	26m
csi-snapshot-controller	4.16.0	True	False	False	37m
dns	4.16.0	True	False	False	37m
etcd	4.16.0	True	False	False	36m
image-registry	4.16.0	True	False	False	31m
ingress	4.16.0	True	False	False	30m
insights	4.16.0	True	False	False	31m
kube-apiserver	4.16.0	True	False	False	26m
kube-controller-manager	4.16.0	True	False	False	36m
kube-scheduler	4.16.0	True	False	False	36m
kube-storage-version-migrator	4.16.0	True	False	False	37m
machine-api	4.16.0	True	False	False	29m
machine-approver	4.16.0	True	False	False	37m
machine-config	4.16.0	True	False	False	36m
marketplace	4.16.0	True	False	False	37m
monitoring	4.16.0	True	False	False	29m
network	4.16.0	True	False	False	38m
node-tuning	4.16.0	True	False	False	37m
openshift-apiserver	4.16.0	True	False	False	32m
openshift-controller-manager	4.16.0	True	False	False	30m
openshift-samples	4.16.0	True	False	False	32m
operator-lifecycle-manager	4.16.0	True	False	False	37m
operator-lifecycle-manager-catalog	4.16.0	True	False	False	37m
operator-lifecycle-manager-packageserver	4.16.0	True	False	False	32m
service-ca	4.16.0	True	False	False	38m
storage	4.16.0	True	False	False	37m

2. Configure the Operators that are not available.

6.4.4.15.1. Disabling the default OperatorHub catalog sources

Operator catalogs that source content provided by Red Hat and community projects are configured for OperatorHub by default during an OpenShift Container Platform installation. In a restricted network environment, you must disable the default catalogs as a cluster administrator.

Procedure

- Disable the sources for the default catalogs by adding **disableAllDefaultSources: true** to the **OperatorHub** object:

```
$ oc patch OperatorHub cluster --type json \
  -p '[{"op": "add", "path": "/spec/disableAllDefaultSources", "value": true}]'
```

TIP

Alternatively, you can use the web console to manage catalog sources. From the **Administration** → **Cluster Settings** → **Configuration** → **OperatorHub** page, click the **Sources** tab, where you can create, update, delete, disable, and enable individual sources.

6.4.4.15.2. Image registry storage configuration

Amazon Web Services provides default storage, which means the Image Registry Operator is available after installation. However, if the Registry Operator cannot create an S3 bucket and automatically configure storage, you must manually configure registry storage.

Instructions are shown for configuring a persistent volume, which is required for production clusters. Where applicable, instructions are shown for configuring an empty directory as the storage location, which is available for only non-production clusters.

Additional instructions are provided for allowing the image registry to use block storage types by using the **Recreate** rollout strategy during upgrades.

6.4.4.15.2.1. Configuring registry storage for AWS with user-provisioned infrastructure

During installation, your cloud credentials are sufficient to create an Amazon S3 bucket and the Registry Operator will automatically configure storage.

If the Registry Operator cannot create an S3 bucket and automatically configure storage, you can create an S3 bucket and configure storage with the following procedure.

Prerequisites

- You have a cluster on AWS with user-provisioned infrastructure.
- For Amazon S3 storage, the secret is expected to contain two keys:
 - **REGISTRY_STORAGE_S3_ACCESSKEY**
 - **REGISTRY_STORAGE_S3_SECRETKEY**

Procedure

Use the following procedure if the Registry Operator cannot create an S3 bucket and automatically configure storage.

1. Set up a [Bucket Lifecycle Policy](#) to abort incomplete multipart uploads that are one day old.
2. Fill in the storage configuration in **configs.imageregistry.operator.openshift.io/cluster**:

```
$ oc edit configs.imageregistry.operator.openshift.io/cluster
```

Example configuration

```
storage:
  s3:
    bucket: <bucket-name>
    region: <region-name>
```



WARNING

To secure your registry images in AWS, [block public access](#) to the S3 bucket.

6.4.4.15.2.2. Configuring storage for the image registry in non-production clusters

You must configure storage for the Image Registry Operator. For non-production clusters, you can set the image registry to an empty directory. If you do so, all images are lost if you restart the registry.

Procedure

- To set the image registry storage to an empty directory:

```
$ oc patch configs.imageregistry.operator.openshift.io cluster --type merge --patch '{"spec": {"storage":{"emptyDir":{}}}'
```



WARNING

Configure this option for only non-production clusters.

If you run this command before the Image Registry Operator initializes its components, the **oc patch** command fails with the following error:

```
Error from server (NotFound): configs.imageregistry.operator.openshift.io "cluster" not found
```

Wait a few minutes and run the command again.

6.4.4.16. Deleting the bootstrap resources

After you complete the initial Operator configuration for the cluster, remove the bootstrap resources from Amazon Web Services (AWS).

Prerequisites

- You completed the initial Operator configuration for your cluster.

Procedure

- Delete the bootstrap resources. If you used the CloudFormation template, [delete its stack](#):

- Delete the stack by using the AWS CLI:

```
$ aws cloudformation delete-stack --stack-name <name> 1
```

1 **<name>** is the name of your bootstrap stack.

- Delete the stack by using the [AWS CloudFormation console](#).

6.4.4.17. Creating the Ingress DNS Records

If you removed the DNS Zone configuration, manually create DNS records that point to the Ingress load balancer. You can create either a wildcard record or specific records. While the following procedure uses A records, you can use other record types that you require, such as CNAME or alias.

Prerequisites

- You deployed an OpenShift Container Platform cluster on Amazon Web Services (AWS) that uses infrastructure that you provisioned.
- You installed the OpenShift CLI (**oc**).
- You installed the **jq** package.
- You downloaded the AWS CLI and installed it on your computer. See [Install the AWS CLI Using the Bundled Installer \(Linux, macOS, or Unix\)](#).

Procedure

- Determine the routes to create.
 - To create a wildcard record, use ***.apps.<cluster_name>.<domain_name>**, where **<cluster_name>** is your cluster name, and **<domain_name>** is the Route 53 base domain for your OpenShift Container Platform cluster.
 - To create specific records, you must create a record for each route that your cluster uses, as shown in the output of the following command:

```
$ oc get --all-namespaces -o jsonpath='{range .items[*]}{range .status.ingress[*]}{.host}\n\n'} routes
```

Example output

```

oauth-openshift.apps.<cluster_name>.<domain_name>
console-openshift-console.apps.<cluster_name>.<domain_name>
downloads-openshift-console.apps.<cluster_name>.<domain_name>
alertmanager-main-openshift-monitoring.apps.<cluster_name>.<domain_name>
prometheus-k8s-openshift-monitoring.apps.<cluster_name>.<domain_name>

```

- Retrieve the Ingress Operator load balancer status and note the value of the external IP address that it uses, which is shown in the **EXTERNAL-IP** column:

```
$ oc -n openshift-ingress get service router-default
```

Example output

```

NAME          TYPE          CLUSTER-IP    EXTERNAL-IP          PORT(S)
AGE
router-default LoadBalancer  172.30.62.215  ab3...28.us-east-2.elb.amazonaws.com
80:31499/TCP,443:30693/TCP  5m

```

- Locate the hosted zone ID for the load balancer:

```
$ aws elb describe-load-balancers | jq -r '.LoadBalancerDescriptions[] | select(.DNSName == "<external_ip>").CanonicalHostedZoneNameID' 1
```

- For **<external_ip>**, specify the value of the external IP address of the Ingress Operator load balancer that you obtained.

Example output

```
Z3AADJGX6KTTL2
```

The output of this command is the load balancer hosted zone ID.

- Obtain the public hosted zone ID for your cluster's domain:

```

$ aws route53 list-hosted-zones-by-name \
  --dns-name "<domain_name>" 1
  --query 'HostedZones[? Config.PrivateZone != `true` && Name ==
`<domain_name>.`].Id' 2
  --output text

```

- For **<domain_name>**, specify the Route 53 base domain for your OpenShift Container Platform cluster.

Example output

```
/hostedzone/Z3URY6TWQ91KVV
```

The public hosted zone ID for your domain is shown in the command output. In this example, it is **Z3URY6TWQ91KVV**.

- Add the alias records to your private zone:

```
$ aws route53 change-resource-record-sets --hosted-zone-id "<private_hosted_zone_id>" --
change-batch '{ ❶
> "Changes": [
> {
>   "Action": "CREATE",
>   "ResourceRecordSet": {
>     "Name": "\\052.apps.<cluster_domain>", ❷
>     "Type": "A",
>     "AliasTarget":{
>       "HostedZoneId": "<hosted_zone_id>", ❸
>       "DNSName": "<external_ip>.", ❹
>       "EvaluateTargetHealth": false
>     }
>   }
> }
> ]
> }'
```

- ❶ For **<private_hosted_zone_id>**, specify the value from the output of the CloudFormation template for DNS and load balancing.
- ❷ For **<cluster_domain>**, specify the domain or subdomain that you use with your OpenShift Container Platform cluster.
- ❸ For **<hosted_zone_id>**, specify the public hosted zone ID for the load balancer that you obtained.
- ❹ For **<external_ip>**, specify the value of the external IP address of the Ingress Operator load balancer. Ensure that you include the trailing period (.) in this parameter value.

6. Add the records to your public zone:

```
$ aws route53 change-resource-record-sets --hosted-zone-id "<public_hosted_zone_id>" --
change-batch '{ ❶
> "Changes": [
> {
>   "Action": "CREATE",
>   "ResourceRecordSet": {
>     "Name": "\\052.apps.<cluster_domain>", ❷
>     "Type": "A",
>     "AliasTarget":{
>       "HostedZoneId": "<hosted_zone_id>", ❸
>       "DNSName": "<external_ip>.", ❹
>       "EvaluateTargetHealth": false
>     }
>   }
> }
> ]
> }'
```

- ❶ For **<public_hosted_zone_id>**, specify the public hosted zone for your domain.
- ❷ For **<cluster_domain>**, specify the domain or subdomain that you use with your OpenShift Container Platform cluster.

- 3 For **<hosted_zone_id>**, specify the public hosted zone ID for the load balancer that you obtained.
- 4 For **<external_ip>**, specify the value of the external IP address of the Ingress Operator load balancer. Ensure that you include the trailing period (.) in this parameter value.

6.4.4.18. Completing an AWS installation on user-provisioned infrastructure

After you start the OpenShift Container Platform installation on Amazon Web Service (AWS) user-provisioned infrastructure, monitor the deployment to completion.

Prerequisites

- You removed the bootstrap node for an OpenShift Container Platform cluster on user-provisioned AWS infrastructure.
- You installed the **oc** CLI.

Procedure


1. From the directory that contains the installation program, complete the cluster installation:

```
$ ./openshift-install --dir <installation_directory> wait-for install-complete 1
```

- 1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

Example output

```
INFO Waiting up to 40m0s for the cluster at https://api.mycluster.example.com:6443 to
initialize...
INFO Waiting up to 10m0s for the openshift-console route to be created...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 1s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

2. Register your cluster on the [Cluster registration](#) page.

6.4.4.19. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

6.4.4.20. Logging in to the cluster by using the web console

The **kubeadmin** user exists by default after an OpenShift Container Platform installation. You can log in to your cluster as the **kubeadmin** user by using the OpenShift Container Platform web console.


Prerequisites

- You have access to the installation host.
- You completed a cluster installation and all cluster Operators are available.

Procedure

1. Obtain the password for the **kubeadmin** user from the **kubeadmin-password** file on the installation host:

```
$ cat <installation_directory>/auth/kubeadmin-password
```



NOTE

Alternatively, you can obtain the **kubeadmin** password from the **<installation_directory>/openshift_install.log** log file on the installation host.

2. List the OpenShift Container Platform web console route:

```
$ oc get routes -n openshift-console | grep 'console-openshift'
```



NOTE

Alternatively, you can obtain the OpenShift Container Platform route from the **<installation_directory>/openshift_install.log** log file on the installation host.

Example output

```
console    console-openshift-console.apps.<cluster_name>.<base_domain>    console
https reencrypt/Redirect None
```

3. Navigate to the route detailed in the output of the preceding command in a web browser and log in as the **kubeadmin** user.

Additional resources

- See [Accessing the web console](#) for more details about accessing and understanding the OpenShift Container Platform web console.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service

6.4.4.21. Additional resources

- See [Working with stacks](#) in the AWS documentation for more information about AWS CloudFormation stacks.

6.4.4.22. Next steps


- [Validate an installation.](#)

- [Customize your cluster](#).
- [Configure image streams](#) for the Cluster Samples Operator and the **must-gather** tool.
- Learn how to [use Operator Lifecycle Manager \(OLM\) on restricted networks](#) .
- If the mirror registry that you used to install your cluster has a trusted CA, add it to the cluster by [configuring additional trust stores](#).
- If necessary, you can [opt out of remote health reporting](#) .
- If necessary, see [Registering your disconnected cluster](#)
- If necessary, you can [remove cloud provider credentials](#).

6.5. INSTALLING A THREE-NODE CLUSTER ON AWS

In OpenShift Container Platform version 4.16, you can install a three-node cluster on Amazon Web Services (AWS). A three-node cluster consists of three control plane machines, which also act as compute machines. This type of cluster provides a smaller, more resource efficient cluster, for cluster administrators and developers to use for testing, development, and production.

You can install a three-node cluster using either installer-provisioned or user-provisioned infrastructure.




NOTE

Deploying a three-node cluster using an AWS Marketplace image is not supported.

6.5.1. Configuring a three-node cluster

You configure a three-node cluster by setting the number of worker nodes to **0** in the **install-config.yaml** file before deploying the cluster. Setting the number of worker nodes to **0** ensures that the control plane machines are schedulable. This allows application workloads to be scheduled to run from the control plane nodes.



NOTE

Because application workloads run from control plane nodes, additional subscriptions are required, as the control plane nodes are considered to be compute nodes.

Prerequisites

- You have an existing **install-config.yaml** file.

Procedure

1. Set the number of compute replicas to **0** in your **install-config.yaml** file, as shown in the following **compute** stanza:

Example **install-config.yaml** file for a three-node cluster

```
apiVersion: v1
baseDomain: example.com
compute:
```



```
- name: worker
  platform: {}
  replicas: 0
# ...
```

2. If you are deploying a cluster with user-provisioned infrastructure:

- After you create the Kubernetes manifest files, make sure that the **spec.mastersSchedulable** parameter is set to **true** in **cluster-scheduler-02-config.yml** file. You can locate this file in **<installation_directory>/manifests**. For more information, see "Creating the Kubernetes manifest and Ignition config files" in "Installing a cluster on user-provisioned infrastructure in AWS by using CloudFormation templates".
- Do not create additional worker nodes.

Example cluster-scheduler-02-config.yml file for a three-node cluster

```
apiVersion: config.openshift.io/v1
kind: Scheduler
metadata:
  creationTimestamp: null
  name: cluster
spec:
  mastersSchedulable: true
policy:
  name: ""
status: {}
```

6.5.2. Next steps

- [Installing a cluster on AWS with customizations](#)
- [Installing a cluster on user-provisioned infrastructure in AWS by using CloudFormation templates](#)

6.6. UNINSTALLING A CLUSTER ON AWS

You can remove a cluster that you deployed to Amazon Web Services (AWS).

6.6.1. Removing a cluster that uses installer-provisioned infrastructure

You can remove a cluster that uses installer-provisioned infrastructure from your cloud.



NOTE

After uninstallation, check your cloud provider for any resources not removed properly, especially with User Provisioned Infrastructure (UPI) clusters. There might be resources that the installer did not create or that the installer is unable to access.

Prerequisites

- You have a copy of the installation program that you used to deploy the cluster.
- You have the files that the installation program generated when you created your cluster.

Procedure

1. From the directory that contains the installation program on the computer that you used to install the cluster, run the following command:

```
$ ./openshift-install destroy cluster \
--dir <installation_directory> --log-level info 1 2
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.
- 2** To view different details, specify **warn**, **debug**, or **error** instead of **info**.



NOTE

You must specify the directory that contains the cluster definition files for your cluster. The installation program requires the **metadata.json** file in this directory to delete the cluster.

2. Optional: Delete the **<installation_directory>** directory and the OpenShift Container Platform installation program.

6.6.2. Deleting Amazon Web Services resources with the Cloud Credential Operator utility

After uninstalling an OpenShift Container Platform cluster that uses short-term credentials managed outside the cluster, you can use the CCO utility (**ccoctl**) to remove the Amazon Web Services (AWS) resources that **ccoctl** created during installation.

Prerequisites

- Extract and prepare the **ccoctl** binary.
- Uninstall an OpenShift Container Platform cluster on AWS that uses short-term credentials.

Procedure

- Delete the AWS resources that **ccoctl** created by running the following command:

```
$ ccoctl aws delete \
--name=<name> 1 \
--region=<aws_region> 2
```

- 1** **<name>** matches the name that was originally used to create and tag the cloud resources.
- 2** **<aws_region>** is the AWS region in which to delete cloud resources.

Example output

```
2021/04/08 17:50:41 Identity Provider object .well-known/openid-configuration deleted from
the bucket <name>-oidc
2021/04/08 17:50:42 Identity Provider object keys.json deleted from the bucket <name>-oidc
```

```

2021/04/08 17:50:43 Identity Provider bucket <name>-oidc deleted
2021/04/08 17:51:05 Policy <name>-openshift-cloud-credential-operator-cloud-credential-o
associated with IAM Role <name>-openshift-cloud-credential-operator-cloud-credential-o
deleted
2021/04/08 17:51:05 IAM Role <name>-openshift-cloud-credential-operator-cloud-credential-
o deleted
2021/04/08 17:51:07 Policy <name>-openshift-cluster-csi-drivers-ebs-cloud-credentials
associated with IAM Role <name>-openshift-cluster-csi-drivers-ebs-cloud-credentials deleted
2021/04/08 17:51:07 IAM Role <name>-openshift-cluster-csi-drivers-ebs-cloud-credentials
deleted
2021/04/08 17:51:08 Policy <name>-openshift-image-registry-installer-cloud-credentials
associated with IAM Role <name>-openshift-image-registry-installer-cloud-credentials
deleted
2021/04/08 17:51:08 IAM Role <name>-openshift-image-registry-installer-cloud-credentials
deleted
2021/04/08 17:51:09 Policy <name>-openshift-ingress-operator-cloud-credentials associated
with IAM Role <name>-openshift-ingress-operator-cloud-credentials deleted
2021/04/08 17:51:10 IAM Role <name>-openshift-ingress-operator-cloud-credentials deleted
2021/04/08 17:51:11 Policy <name>-openshift-machine-api-aws-cloud-credentials associated
with IAM Role <name>-openshift-machine-api-aws-cloud-credentials deleted
2021/04/08 17:51:11 IAM Role <name>-openshift-machine-api-aws-cloud-credentials deleted
2021/04/08 17:51:39 Identity Provider with ARN arn:aws:iam::<aws_account_id>:oidc-
provider/<name>-oidc.s3.<aws_region>.amazonaws.com deleted

```

Verification

- To verify that the resources are deleted, query AWS. For more information, refer to AWS documentation.

6.6.3. Deleting a cluster with a configured AWS Local Zone infrastructure

After you install a cluster on Amazon Web Services (AWS) into an existing Virtual Private Cloud (VPC), and you set subnets for each Local Zone location, you can delete the cluster and any AWS resources associated with it.

The example in the procedure assumes that you created a VPC and its subnets by using a CloudFormation template.

Prerequisites

- You know the name of the CloudFormation stacks, **<local_zone_stack_name>** and **<vpc_stack_name>**, that were used during the creation of the network. You need the name of the stack to delete the cluster.
- You have access rights to the directory that contains the installation files that were created by the installation program.
- Your account includes a policy that provides you with permissions to delete the CloudFormation stack.

Procedure

1. Change to the directory that contains the stored installation program, and delete the cluster by using the **destroy cluster** command:

```

$ ./openshift-install destroy cluster --dir <installation_directory> \ 1
--log-level=debug 2

```

- 1** For **<installation_directory>**, specify the directory that stored any files created by the installation program.
- 2** To view different log details, specify **error**, **info**, or **warn** instead of **debug**.

2. Delete the CloudFormation stack for the Local Zone subnet:

```
$ aws cloudformation delete-stack --stack-name <local_zone_stack_name>
```

3. Delete the stack of resources that represent the VPC:

```
$ aws cloudformation delete-stack --stack-name <vpc_stack_name>
```

Verification

- Check that you removed the stack resources by issuing the following commands in the AWS CLI. The AWS CLI outputs that no template component exists.

```
$ aws cloudformation describe-stacks --stack-name <local_zone_stack_name>
```

```
$ aws cloudformation describe-stacks --stack-name <vpc_stack_name>
```

Additional resources

- See [Working with stacks](#) in the AWS documentation for more information about AWS CloudFormation stacks.
- [Opt into AWS Local Zones](#)
- [AWS Local Zones available locations](#)
- [AWS Local Zones features](#)

6.7. INSTALLATION CONFIGURATION PARAMETERS FOR AWS

Before you deploy an OpenShift Container Platform cluster on AWS, you provide parameters to customize your cluster and the platform that hosts it. When you create the **install-config.yaml** file, you provide values for the required parameters through the command line. You can then modify the **install-config.yaml** file to customize your cluster further.

6.7.1. Available installation configuration parameters for AWS

The following tables specify the required, optional, and AWS-specific installation configuration parameters that you can set as part of the installation process.



NOTE

After installation, you cannot modify these parameters in the **install-config.yaml** file.

6.7.1.1. Required configuration parameters

Required installation configuration parameters are described in the following table:

Table 6.27. Required parameters


Parameter	Description	Values
<code>apiVersion:</code>	The API version for the install-config.yaml content. The current version is v1 . The installation program may also support older API versions.	String
<code>baseDomain:</code>	The base domain of your cloud provider. The base domain is used to create routes to your OpenShift Container Platform cluster components. The full DNS name for your cluster is a combination of the baseDomain and metadata.name parameter values that uses the <metadata.name> . <baseDomain> format.	A fully-qualified domain or subdomain name, such as example.com .
<code>metadata:</code>	Kubernetes resource ObjectMeta , from which only the name parameter is consumed.	Object
<code>metadata: name:</code>	The name of the cluster. DNS records for the cluster are all subdomains of {{.metadata.name}} . {{.baseDomain}} .	String of lowercase letters, hyphens (-), and periods (.), such as dev .
<code>platform:</code>	The configuration for the specific platform upon which to perform the installation: aws , baremetal , azure , gcp , ibmcloud , nutanix , openstack , powervs , vsphere , or {} . For additional information about platform . <platform> parameters, consult the table for your specific platform that follows.	Object

Parameter	Description	Values
<code>pullSecret:</code>	Get a pull secret from Red Hat OpenShift Cluster Manager to authenticate downloading container images for OpenShift Container Platform components from services such as Quay.io.	<pre>{ "auths":{ "cloud.openshift.com":{ "auth":"b3Blb=", "email":"you@example.com" }, "quay.io":{ "auth":"b3Blb=", "email":"you@example.com" } } }</pre>

6.7.1.2. Network configuration parameters

You can customize your installation configuration based on the requirements of your existing network infrastructure. For example, you can expand the IP address block for the cluster network or provide different IP address blocks than the defaults.

Only IPv4 addresses are supported.




NOTE

Globalnet is not supported with Red Hat OpenShift Data Foundation disaster recovery solutions. For regional disaster recovery scenarios, ensure that you use a nonoverlapping range of private IP addresses for the cluster and service networks in each cluster.

Table 6.28. Network parameters

Parameter	Description	Values
<code>networking:</code>	The configuration for the cluster network.	Object <div style="display: flex; align-items: center; margin-top: 10px;"> <div> <p>NOTE</p> <p>You cannot modify parameters specified by the networking object after installation.</p> </div> </div>
<code>networking: networkType:</code>	The Red Hat OpenShift Networking network plugin to install.	OVNKubernetes. OVNKubernetes is a CNI plugin for Linux networks and hybrid networks that contain both Linux and Windows servers. The default value is OVNKubernetes .


Parameter	Description	Values
<code>networking: clusterNetwork:</code>	<p>The IP address blocks for pods.</p> <p>The default value is 10.128.0.0/14 with a host prefix of /23.</p> <p>If you specify multiple IP address blocks, the blocks must not overlap.</p>	<p>An array of objects. For example:</p> <pre>networking: clusterNetwork: - cidr: 10.128.0.0/14 hostPrefix: 23</pre>
<code>networking: clusterNetwork: cidr:</code>	<p>Required if you use networking.clusterNetwork. An IP address block.</p> <p>An IPv4 network.</p>	<p>An IP address block in Classless Inter-Domain Routing (CIDR) notation. The prefix length for an IPv4 block is between 0 and 32.</p>
<code>networking: clusterNetwork: hostPrefix:</code>	<p>The subnet prefix length to assign to each individual node. For example, if hostPrefix is set to 23 then each node is assigned a /23 subnet out of the given cidr. A hostPrefix value of 23 provides $510 (2^{(32 - 23)} - 2)$ pod IP addresses.</p>	<p>A subnet prefix.</p> <p>The default value is 23.</p>
<code>networking: serviceNetwork:</code>	<p>The IP address block for services. The default value is 172.30.0.0/16.</p> <p>The OVN-Kubernetes network plugins supports only a single IP address block for the service network.</p>	<p>An array with an IP address block in CIDR format. For example:</p> <pre>networking: serviceNetwork: - 172.30.0.0/16</pre>
<code>networking: machineNetwork:</code>	<p>The IP address blocks for machines.</p> <p>If you specify multiple IP address blocks, the blocks must not overlap.</p>	<p>An array of objects. For example:</p> <pre>networking: machineNetwork: - cidr: 10.0.0.0/16</pre>
<code>networking: machineNetwork: cidr:</code>	<p>Required if you use networking.machineNetwork. An IP address block. The default value is 10.0.0.0/16 for all platforms other than libvirt and IBM Power® Virtual Server. For libvirt, the default value is 192.168.126.0/24. For IBM Power® Virtual Server, the default value is 192.168.0.0/24.</p>	<p>An IP network block in CIDR notation.</p> <p>For example, 10.0.0.0/16.</p> <div style="display: flex; align-items: flex-start;">  <div> <p>NOTE</p> <p>Set the networking.machineNetwork to match the CIDR that the preferred NIC resides in.</p> </div> </div>


6.7.1.3. Optional configuration parameters

Optional installation configuration parameters are described in the following table:

Table 6.29. Optional parameters

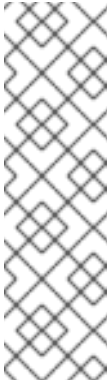
Parameter	Description	Values
<code>additionalTrustBundle:</code>	A PEM-encoded X.509 certificate bundle that is added to the nodes' trusted certificate store. This trust bundle may also be used when a proxy has been configured.	String
<code>capabilities:</code>	Controls the installation of optional core cluster components. You can reduce the footprint of your OpenShift Container Platform cluster by disabling optional components. For more information, see the "Cluster capabilities" page in <i>Installing</i> .	String array
<code>capabilities: baselineCapabilitySet:</code>	Selects an initial set of optional capabilities to enable. Valid values are None , v4.11 , v4.12 and vCurrent . The default value is vCurrent .	String
<code>capabilities: additionalEnabledCapabilities:</code>	Extends the set of optional capabilities beyond what you specify in baselineCapabilitySet . You may specify multiple capabilities in this parameter.	String array
<code>cpuPartitioningMode:</code>	Enables workload partitioning, which isolates OpenShift Container Platform services, cluster management workloads, and infrastructure pods to run on a reserved set of CPUs. Workload partitioning can only be enabled during installation and cannot be disabled after installation. While this field enables workload partitioning, it does not configure workloads to use specific CPUs. For more information, see the <i>Workload partitioning</i> page in the <i>Scalability and Performance</i> section.	None or AllNodes . None is the default value.
<code>compute:</code>	The configuration for the machines that comprise the compute nodes.	Array of MachinePool objects.

Parameter	Description	Values
<code>compute: architecture:</code>	Determines the instruction set architecture of the machines in the pool. Currently, clusters with varied architectures are not supported. All pools must specify the same architecture. Valid values are amd64 and arm64 . Not all installation options support the 64-bit ARM architecture. To verify if your installation option is supported on your platform, see <i>Supported installation methods for different platforms</i> in <i>Selecting a cluster installation method and preparing it for users</i> .	String
<code>compute: hyperthreading:</code>	Whether to enable or disable simultaneous multithreading, or hyperthreading , on compute machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.  IMPORTANT If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.	Enabled or Disabled
<code>compute: name:</code>	Required if you use compute . The name of the machine pool.	worker
<code>compute: platform:</code>	Required if you use compute . Use this parameter to specify the cloud provider to host the worker machines. This parameter value must match the controlPlane.platform parameter value.	aws, azure, gcp, ibmcloud, nutanix, openstack, powervs, vsphere, or {}
<code>compute: replicas:</code>	The number of compute machines, which are also known as worker machines, to provision.	A positive integer greater than or equal to 2 . The default value is 3 .

Parameter	Description	Values
<code>featureSet:</code>	Enables the cluster for a feature set. A feature set is a collection of OpenShift Container Platform features that are not enabled by default. For more information about enabling a feature set during installation, see "Enabling features using feature gates".	String. The name of the feature set to enable, such as TechPreviewNoUpgrade .
<code>controlPlane:</code>	The configuration for the machines that comprise the control plane.	Array of MachinePool objects.
<code>controlPlane: architecture:</code>	Determines the instruction set architecture of the machines in the pool. Currently, clusters with varied architectures are not supported. All pools must specify the same architecture. Valid values are amd64 and arm64 . Not all installation options support the 64-bit ARM architecture. To verify if your installation option is supported on your platform, see <i>Supported installation methods for different platforms</i> in <i>Selecting a cluster installation method and preparing it for users</i> .	String
<code>controlPlane: hyperthreading:</code>	Whether to enable or disable simultaneous multithreading, or hyperthreading , on control plane machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.  IMPORTANT If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.	Enabled or Disabled

Parameter	Description	Values
controlPlane: name:	Required if you use controlPlane . The name of the machine pool.	master
controlPlane: platform:	Required if you use controlPlane . Use this parameter to specify the cloud provider that hosts the control plane machines. This parameter value must match the compute.platform parameter value.	aws, azure, gcp, ibmcloud, nutanix, openstack, powervs, vsphere, or {}
controlPlane: replicas:	The number of control plane machines to provision.	Supported values are 3 , or 1 when deploying single-node OpenShift.
credentialsMode:	The Cloud Credential Operator (CCO) mode. If no mode is specified, the CCO dynamically tries to determine the capabilities of the provided credentials, with a preference for mint mode on the platforms where multiple modes are supported.	Mint, Passthrough, Manual or an empty string (""). ^[1]

Parameter	Description	Values
<p>fips:</p>	<p>Enable or disable FIPS mode. The default is false (disabled). If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.</p> <div style="display: flex; align-items: flex-start;"> <div style="width: 20px; height: 100%; background-color: black; margin-right: 10px;"></div> <div> <p>IMPORTANT</p> <p>To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see Installing the system in FIPS mode.</p> <p>When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.</p> </div> </div> <div style="display: flex; align-items: flex-start; margin-top: 20px;"> <div style="width: 20px; height: 100%; background-color: black; margin-right: 10px;"></div> <div> <p>NOTE</p> <p>If you are using Azure File storage, you cannot enable FIPS mode.</p> </div> </div>	<p>false or true</p>
<p>imageContentSources:</p>	<p>Sources and repositories for the release-image content.</p>	<p>Array of objects. Includes a source and, optionally, mirrors, as described in the following rows of this table.</p>

Parameter	Description	Values
<code>imageContentSources:</code> <code>source:</code>	Required if you use imageContentSources . Specify the repository that users refer to, for example, in image pull specifications.	String
<code>imageContentSources:</code> <code>mirrors:</code>	Specify one or more repositories that may also contain the same images.	Array of strings
<code>platform:</code> <code>aws:</code> <code>lbType:</code>	Required to set the NLB load balancer type in AWS. Valid values are Classic or NLB . If no value is specified, the installation program defaults to Classic . The installation program sets the value provided here in the ingress cluster configuration object. If you do not specify a load balancer type for other Ingress Controllers, they use the type set in this parameter.	Classic or NLB . The default value is Classic .
<code>publish:</code>	How to publish or expose the user-facing endpoints of your cluster, such as the Kubernetes API, OpenShift routes.	Internal or External . To deploy a private cluster, which cannot be accessed from the internet, set publish to Internal . The default value is External .
<code>sshKey:</code>	The SSH key to authenticate access to your cluster machines.  NOTE For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your ssh-agent process uses.	For example, sshKey: ssh-ed25519 AAAA...

1. Not all CCO modes are supported for all cloud providers. For more information about CCO modes, see the "Managing cloud provider credentials" entry in the *Authentication and authorization* content.

**NOTE**

If your AWS account has service control policies (SCP) enabled, you must configure the **credentialsMode** parameter to **Mint**, **Passthrough**, or **Manual**.

**IMPORTANT**

Setting this parameter to **Manual** enables alternatives to storing administrator-level secrets in the **kube-system** project, which require additional configuration steps. For more information, see "Alternatives to storing administrator-level secrets in the kube-system project".


6.7.1.4. Optional AWS configuration parameters

Optional AWS configuration parameters are described in the following table:

Table 6.30. Optional AWS parameters


Parameter	Description	Values
compute: platform: aws: amiID:	The AWS AMI used to boot compute machines for the cluster. This is required for regions that require a custom RHCOS AMI.	Any published or custom RHCOS AMI that belongs to the set AWS region. See <i>RHCOS AMIs for AWS infrastructure</i> for available AMI IDs.
compute: platform: aws: iamRole:	A pre-existing AWS IAM role applied to the compute machine pool instance profiles. You can use these fields to match naming schemes and include predefined permissions boundaries for your IAM roles. If undefined, the installation program creates a new IAM role.	The name of a valid AWS IAM role.
compute: platform: aws: rootVolume: iops:	The Input/Output Operations Per Second (IOPS) that is reserved for the root volume.	Integer, for example 4000 .


Parameter	Description	Values
<pre>compute: platform: aws: rootVolume: size:</pre>	The size in GiB of the root volume.	Integer, for example 500 .
<pre>compute: platform: aws: rootVolume: type:</pre>	The type of the root volume.	Valid AWS EBS volume type , such as io1 .
<pre>compute: platform: aws: rootVolume: kmsKeyARN:</pre>	The Amazon Resource Name (key ARN) of a KMS key. This is required to encrypt operating system volumes of worker nodes with a specific KMS key.	Valid key ID or the key ARN
<pre>compute: platform: aws: type:</pre>	The EC2 instance type for the compute machines.	Valid AWS instance type, such as m4.2xlarge . See the Supported AWS machine types table that follows.
<pre>compute: platform: aws: zones:</pre>	The availability zones where the installation program creates machines for the compute machine pool. If you provide your own VPC, you must provide a subnet in that availability zone.	A list of valid AWS availability zones, such as us-east-1c , in a YAML sequence .

Parameter	Description	Values
<pre>compute: aws: region:</pre>	The AWS region that the installation program creates compute resources in.	<p>Any valid AWS region, such as us-east-1. You can use the AWS CLI to access the regions available based on your selected instance type. For example:</p> <pre>aws ec2 describe-instance-type-offerings -- filters Name=instance- type,Values=c7g.xlarge</pre> <div style="display: flex; align-items: flex-start;">  <div> <p>IMPORTANT</p> <p>When running on ARM based AWS instances, ensure that you enter a region where AWS Graviton processors are available. See Global availability map in the AWS documentation. Currently, AWS Graviton3 processors are only available in some regions.</p> </div> </div>
<pre>controlPlane: platform: aws: amiID:</pre>	The AWS AMI used to boot control plane machines for the cluster. This is required for regions that require a custom RHCOS AMI.	Any published or custom RHCOS AMI that belongs to the set AWS region. See <i>RHCOS AMIs for AWS infrastructure</i> for available AMI IDs.
<pre>controlPlane: platform: aws: iamRole:</pre>	A pre-existing AWS IAM role applied to the control plane machine pool instance profiles. You can use these fields to match naming schemes and include predefined permissions boundaries for your IAM roles. If undefined, the installation program creates a new IAM role.	The name of a valid AWS IAM role.
<pre>controlPlane: platform: aws: rootVolume: iops:</pre>	The Input/Output Operations Per Second (IOPS) that is reserved for the root volume on control plane machines.	Integer, for example 4000 .

Parameter	Description	Values
<pre>controlPlane: platform: aws: rootVolume: size:</pre>	The size in GiB of the root volume for control plane machines.	Integer, for example 500 .
<pre>controlPlane: platform: aws: rootVolume: type:</pre>	The type of the root volume for control plane machines.	Valid AWS EBS volume type , such as io1 .
<pre>controlPlane: platform: aws: rootVolume: kmsKeyARN:</pre>	The Amazon Resource Name (key ARN) of a KMS key. This is required to encrypt operating system volumes of control plane nodes with a specific KMS key.	Valid key ID and the key ARN
<pre>controlPlane: platform: aws: type:</pre>	The EC2 instance type for the control plane machines.	Valid AWS instance type, such as m6i.xlarge . See the Supported AWS machine types table that follows.
<pre>controlPlane: platform: aws: zones:</pre>	The availability zones where the installation program creates machines for the control plane machine pool.	A list of valid AWS availability zones, such as us-east-1c , in a YAML sequence .
<pre>controlPlane: aws: region:</pre>	The AWS region that the installation program creates control plane resources in.	Valid AWS region , such as us-east-1 .

Parameter	Description	Values
platform: aws: amiID:	The AWS AMI used to boot all machines for the cluster. If set, the AMI must belong to the same region as the cluster. This is required for regions that require a custom RHCOS AMI.	Any published or custom RHCOS AMI that belongs to the set AWS region. See <i>RHCOS AMIs for AWS infrastructure</i> for available AMI IDs.
platform: aws: hostedZone:	An existing Route 53 private hosted zone for the cluster. You can only use a pre-existing hosted zone when also supplying your own VPC. The hosted zone must already be associated with the user-provided VPC before installation. Also, the domain of the hosted zone must be the cluster domain or a parent of the cluster domain. If undefined, the installation program creates a new hosted zone.	String, for example Z3URY6TWQ91KVV .
platform: aws: hostedZoneRole:	An Amazon Resource Name (ARN) for an existing IAM role in the account containing the specified hosted zone. The installation program and cluster operators will assume this role when performing operations on the hosted zone. This parameter should only be used if you are installing a cluster into a shared VPC.	String, for example arn:aws:iam::1234567890:role/shared-vpc-role .
platform: aws: serviceEndpoints: - name: url:	The AWS service endpoint name and URL. Custom endpoints are only required for cases where alternative AWS endpoints, like FIPS, must be used. Custom API endpoints can be specified for EC2, S3, IAM, Elastic Load Balancing, Tagging, Route 53, and STS AWS services.	Valid AWS service endpoint name and valid AWS service endpoint URL.

Parameter	Description	Values
<pre>platform: aws: userTags:</pre>	<p>A map of keys and values that the installation program adds as tags to all resources that it creates.</p>	<p>Any valid YAML map, such as key value pairs in the <key>: <value> format. For more information about AWS tags, see Tagging Your Amazon EC2 Resources in the AWS documentation.</p> <div style="display: flex; align-items: flex-start;">  <div> <p>NOTE</p> <p>You can add up to 25 user defined tags during installation. The remaining 25 tags are reserved for OpenShift Container Platform.</p> </div> </div>
<pre>platform: aws: propagateUserTags:</pre>	<p>A flag that directs in-cluster Operators to include the specified user tags in the tags of the AWS resources that the Operators create.</p>	<p>Boolean values, for example true or false.</p>
<pre>platform: aws: subnets:</pre>	<p>If you provide the VPC instead of allowing the installation program to create the VPC for you, specify the subnet for the cluster to use. The subnet must be part of the same machineNetwork[].cidr ranges that you specify.</p> <p>For a standard cluster, specify a public and a private subnet for each availability zone.</p> <p>For a private cluster, specify a private subnet for each availability zone.</p> <p>For clusters that use AWS Local Zones, you must add AWS Local Zone subnets to this list to ensure edge machine pool creation.</p>	<p>Valid subnet IDs.</p>

Parameter	Description	Values
<p>platform: aws:</p> <p>publicIpv4Pool :</p>	<p>The public IPv4 pool ID that is used to allocate Elastic IPs (EIPs) when publish is set to External. You must provision and advertise the pool in the same AWS account and region of the cluster. You must ensure that you have $2n + 1$ IPv4 available in the pool where n is the total number of AWS zones used to deploy the Network Load Balancer (NLB) for API, NAT gateways, and bootstrap node. For more information about bring your own IP addresses (BYOIP) in AWS, see Onboard your BYOIP.</p>	<p>A valid public IPv4 pool id</p>  <p>NOTE</p> <p>BYOIP can be enabled only for customized installations that have no network restrictions.</p>
<p>platform: aws:</p> <p>preserveBoots trapIgnition:</p>	<p>Prevents the S3 bucket from being deleted after completion of bootstrapping.</p>	<p>true or false. The default value is false, which results in the S3 bucket being deleted.</p>

CHAPTER 7. INSTALLING ON AZURE

7.1. PREPARING TO INSTALL ON AZURE

7.1.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).

7.1.2. Requirements for installing OpenShift Container Platform on Azure

Before installing OpenShift Container Platform on Microsoft Azure, you must configure an Azure account. See [Configuring an Azure account](#) for details about account configuration, account limits, public DNS zone configuration, required roles, creating service principals, and supported Azure regions.

If the cloud identity and access management (IAM) APIs are not accessible in your environment, or if you do not want to store an administrator-level credential secret in the **kube-system** namespace, see [Alternatives to storing administrator-level secrets in the kube-system project](#) for other options.

7.1.3. Choosing a method to install OpenShift Container Platform on Azure

You can install OpenShift Container Platform on installer-provisioned or user-provisioned infrastructure. The default installation type uses installer-provisioned infrastructure, where the installation program provisions the underlying infrastructure for the cluster. You can also install OpenShift Container Platform on infrastructure that you provision. If you do not use infrastructure that the installation program provisions, you must manage and maintain the cluster resources yourself.

See [Installation process](#) for more information about installer-provisioned and user-provisioned installation processes.

7.1.3.1. Installing a cluster on installer-provisioned infrastructure

You can install a cluster on Azure infrastructure that is provisioned by the OpenShift Container Platform installation program, by using one of the following methods:

- **Installing a cluster quickly on Azure** You can install OpenShift Container Platform on Azure infrastructure that is provisioned by the OpenShift Container Platform installation program. You can install a cluster quickly by using the default configuration options.
- **Installing a customized cluster on Azure** You can install a customized cluster on Azure infrastructure that the installation program provisions. The installation program allows for some customization to be applied at the installation stage. Many other customization options are available [post-installation](#).
- **Installing a cluster on Azure with network customizations** You can customize your OpenShift Container Platform network configuration during installation, so that your cluster can coexist with your existing IP address allocations and adhere to your network requirements.
- **Installing a cluster on Azure into an existing VNet** You can install OpenShift Container Platform on an existing Azure Virtual Network (VNet) on Azure. You can use this installation method if you have constraints set by the guidelines of your company, such as limits when

creating new accounts or infrastructure.

- **Installing a private cluster on Azure** You can install a private cluster into an existing Azure Virtual Network (VNet) on Azure. You can use this method to deploy OpenShift Container Platform on an internal network that is not visible to the internet.
- **Installing a cluster on Azure into a government region** OpenShift Container Platform can be deployed into Microsoft Azure Government (MAG) regions that are specifically designed for US government agencies at the federal, state, and local level, as well as contractors, educational institutions, and other US customers that must run sensitive workloads on Azure.

7.1.3.2. Installing a cluster on user-provisioned infrastructure

You can install a cluster on Azure infrastructure that you provision, by using the following method:

- **Installing a cluster on Azure using ARM templates** You can install OpenShift Container Platform on Azure by using infrastructure that you provide. You can use the provided Azure Resource Manager (ARM) templates to assist with an installation.

7.1.4. Next steps

- [Configuring an Azure account](#)

7.2. CONFIGURING AN AZURE ACCOUNT

Before you can install OpenShift Container Platform, you must configure a Microsoft Azure account to meet installation requirements.



IMPORTANT

All Azure resources that are available through public endpoints are subject to resource name restrictions, and you cannot create resources that use certain terms. For a list of terms that Azure restricts, see [Resolve reserved resource name errors](#) in the Azure documentation.

7.2.1. Azure account limits

The OpenShift Container Platform cluster uses a number of Microsoft Azure components, and the default [Azure subscription and service limits, quotas, and constraints](#) affect your ability to install OpenShift Container Platform clusters.



IMPORTANT


Default limits vary by offer category types, such as Free Trial and Pay-As-You-Go, and by series, such as Dv2, F, and G. For example, the default for Enterprise Agreement subscriptions is 350 cores.

Check the limits for your subscription type and if necessary, increase quota limits for your account before you install a default cluster on Azure.

The following table summarizes the Azure components whose limits can impact your ability to install and run OpenShift Container Platform clusters.

Component	Number of components required by default	Default Azure limit	Description
vCPU	44	20 per region	<p>A default cluster requires 44 vCPUs, so you must increase the account limit.</p> <p>By default, each cluster creates the following instances:</p> <ul style="list-style-type: none"> • One bootstrap machine, which is removed after installation • Three control plane machines • Three compute machines <p>Because the bootstrap and control plane machines use Standard_D8s_v3 virtual machines, which use 8 vCPUs, and the compute machines use Standard_D4s_v3 virtual machines, which use 4 vCPUs, a default cluster requires 44 vCPUs. The bootstrap node VM, which uses 8 vCPUs, is used only during installation.</p> <p>To deploy more worker nodes, enable autoscaling, deploy large workloads, or use a different instance type, you must further increase the vCPU limit for your account to ensure that your cluster can deploy the machines that you require.</p>
OS Disk	7		<p>Each cluster machine must have a minimum of 100 GB of storage and 300 IOPS. While these are the minimum supported values, faster storage is recommended for production clusters and clusters with intensive workloads. For more information about optimizing storage for performance, see the page titled "Optimizing storage" in the "Scalability and performance" section.</p>
VNet	1	1000 per region	<p>Each default cluster requires one Virtual Network (VNet), which contains two subnets.</p>
Network interfaces	7	65,536 per region	<p>Each default cluster requires seven network interfaces. If you create more machines or your deployed workloads create load balancers, your cluster uses more network interfaces.</p>

Component	Number of components required by default	Default Azure limit	Description						
Network security groups	2	5000	<p>Each cluster creates network security groups for each subnet in the VNet. The default cluster creates network security groups for the control plane and for the compute node subnets:</p> <table border="1"> <tr> <td>control plane</td> <td>Allows the control plane machines to be reached on port 6443 from anywhere</td> </tr> <tr> <td>node</td> <td>Allows worker nodes to be reached from the internet on ports 80 and 443</td> </tr> </table>	control plane	Allows the control plane machines to be reached on port 6443 from anywhere	node	Allows worker nodes to be reached from the internet on ports 80 and 443		
control plane	Allows the control plane machines to be reached on port 6443 from anywhere								
node	Allows worker nodes to be reached from the internet on ports 80 and 443								
Network load balancers	3	1000 per region	<p>Each cluster creates the following load balancers:</p> <table border="1"> <tr> <td>default</td> <td>Public IP address that load balances requests to ports 80 and 443 across worker machines</td> </tr> <tr> <td>internal</td> <td>Private IP address that load balances requests to ports 6443 and 22623 across control plane machines</td> </tr> <tr> <td>external</td> <td>Public IP address that load balances requests to port 6443 across control plane machines</td> </tr> </table> <p>If your applications create more Kubernetes LoadBalancer service objects, your cluster uses more load balancers.</p>	default	Public IP address that load balances requests to ports 80 and 443 across worker machines	internal	Private IP address that load balances requests to ports 6443 and 22623 across control plane machines	external	Public IP address that load balances requests to port 6443 across control plane machines
default	Public IP address that load balances requests to ports 80 and 443 across worker machines								
internal	Private IP address that load balances requests to ports 6443 and 22623 across control plane machines								
external	Public IP address that load balances requests to port 6443 across control plane machines								
Public IP addresses	3		Each of the two public load balancers uses a public IP address. The bootstrap machine also uses a public IP address so that you can SSH into the machine to troubleshoot issues during installation. The IP address for the bootstrap node is used only during installation.						
Private IP addresses	7		The internal load balancer, each of the three control plane machines, and each of the three worker machines each use a private IP address.						

Component	Number of components required by default	Default Azure limit	Description
Spot VM vCPUs (optional)	0 If you configure spot VMs, your cluster must have two spot VM vCPUs for every compute node.	20 per region	<p>This is an optional component. To use spot VMs, you must increase the Azure default limit to at least twice the number of compute nodes in your cluster.</p> <div style="display: flex; align-items: center;">  <div> <p>NOTE</p> <p>Using spot VMs for control plane nodes is not recommended.</p> </div> </div>

Additional resources

- [Optimizing storage](#).

7.2.2. Configuring a public DNS zone in Azure

To install OpenShift Container Platform, the Microsoft Azure account you use must have a dedicated public hosted DNS zone in your account. This zone must be authoritative for the domain. This service provides cluster DNS resolution and name lookup for external connections to the cluster.

Procedure

1. Identify your domain, or subdomain, and registrar. You can transfer an existing domain and registrar or obtain a new one through Azure or another source.



NOTE

For more information about purchasing domains through Azure, see [Buy a custom domain name for Azure App Service](#) in the Azure documentation.

2. If you are using an existing domain and registrar, migrate its DNS to Azure. See [Migrate an active DNS name to Azure App Service](#) in the Azure documentation.
3. Configure DNS for your domain. Follow the steps in the [Tutorial: Host your domain in Azure DNS](#) in the Azure documentation to create a public hosted zone for your domain or subdomain, extract the new authoritative name servers, and update the registrar records for the name servers that your domain uses.
Use an appropriate root domain, such as **openshiftcorp.com**, or subdomain, such as **clusters.openshiftcorp.com**.
4. If you use a subdomain, follow your company's procedures to add its delegation records to the parent domain.

7.2.3. Increasing Azure account limits

To increase an account limit, file a support request on the Azure portal.

**NOTE**

You can increase only one type of quota per support request.

Procedure

1. From the Azure portal, click **Help + support** in the lower left corner.
2. Click **New support request** and then select the required values:
 - a. From the **Issue type** list, select **Service and subscription limits (quotas)**
 - b. From the **Subscription** list, select the subscription to modify.
 - c. From the **Quota type** list, select the quota to increase. For example, select **Compute-VM (cores-vCPUs) subscription limit increases** to increase the number of vCPUs, which is required to install a cluster.
 - d. Click **Next: Solutions**.
3. On the **Problem Details** page, provide the required information for your quota increase:
 - a. Click **Provide details** and provide the required details in the **Quota details** window.
 - b. In the SUPPORT METHOD and CONTACT INFO sections, provide the issue severity and your contact details.
4. Click **Next: Review + create** and then click **Create**.

7.2.4. Recording the subscription and tenant IDs

The installation program requires the subscription and tenant IDs that are associated with your Azure account. You can use the Azure CLI to gather this information.

Prerequisites

- You have installed or updated the [Azure CLI](#).

Procedure

1. Log in to the Azure CLI by running the following command:

```
$ az login
```

2. Ensure that you are using the right subscription:
 - a. View a list of available subscriptions by running the following command:

```
$ az account list --refresh
```

Example output

```
[  
  {  
    "cloudName": "AzureCloud",
```

```

    "id": "8xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
    "isDefault": true,
    "name": "Subscription Name 1",
    "state": "Enabled",
    "tenantId": "6xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
    "user": {
      "name": "you@example.com",
      "type": "user"
    }
  },
  {
    "cloudName": "AzureCloud",
    "id": "9xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
    "isDefault": false,
    "name": "Subscription Name 2",
    "state": "Enabled",
    "tenantId": "7xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
    "user": {
      "name": "you2@example.com",
      "type": "user"
    }
  }
]

```

- b. View the details of the active account, and confirm that this is the subscription you want to use, by running the following command:

```
$ az account show
```

Example output

```

{
  "environmentName": "AzureCloud",
  "id": "8xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
  "isDefault": true,
  "name": "Subscription Name 1",
  "state": "Enabled",
  "tenantId": "6xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
  "user": {
    "name": "you@example.com",
    "type": "user"
  }
}

```

3. If you are not using the right subscription:
 - a. Change the active subscription by running the following command:

```
$ az account set -s <subscription_id>
```

- b. Verify that you are using the subscription you need by running the following command:

```
$ az account show
```

Example output

```
{
  "environmentName": "AzureCloud",
  "id": "9xxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
  "isDefault": true,
  "name": "Subscription Name 2",
  "state": "Enabled",
  "tenantId": "7xxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
  "user": {
    "name": "you2@example.com",
    "type": "user"
  }
}
```

4. Record the **id** and **tenantId** parameter values from the output. You require these values to install an OpenShift Container Platform cluster.

7.2.5. Supported identities to access Azure resources

An OpenShift Container Platform cluster requires an Azure identity to create and manage Azure resources. As such, you need one of the following types of identities to complete the installation:

- A service principal
- A system-assigned managed identity
- A user-assigned managed identity

7.2.5.1. Required Azure roles

An OpenShift Container Platform cluster requires an Azure identity to create and manage Azure resources. Before you create the identity, verify that your environment meets the following requirements:

- The Azure account that you use to create the identity is assigned the **User Access Administrator** and **Contributor** roles. These roles are required when:
 - Creating a service principal or user-assigned managed identity.
 - Enabling a system-assigned managed identity on a virtual machine.
- If you are going to use a service principal to complete the installation, verify that the Azure account that you use to create the identity is assigned the **microsoft.directory/servicePrincipals/createAsOwner** permission in Microsoft Entra ID.

To set roles on the Azure portal, see the [Manage access to Azure resources using RBAC and the Azure portal](#) in the Azure documentation.

7.2.5.2. Required Azure permissions for installer-provisioned infrastructure

The installation program requires access to an Azure service principal or managed identity with the necessary permissions to deploy the cluster and to maintain its daily operation. These permissions must be granted to the Azure subscription that is associated with the identity.

The following options are available to you:

- You can assign the identity the **Contributor** and **User Access Administrator** roles. Assigning these roles is the quickest way to grant all of the required permissions. For more information about assigning roles, see the Azure documentation for [managing access to Azure resources using the Azure portal](#).
- If your organization's security policies require a more restrictive set of permissions, you can create a [custom role](#) with the necessary permissions.

The following permissions are required for creating an OpenShift Container Platform cluster on Microsoft Azure.

Example 7.1. Required permissions for creating authorization resources

- **Microsoft.Authorization/policies/audit/action**
- **Microsoft.Authorization/policies/auditIfNotExists/action**
- **Microsoft.Authorization/roleAssignments/read**
- **Microsoft.Authorization/roleAssignments/write**

Example 7.2. Required permissions for creating compute resources

- **Microsoft.Compute/availabilitySets/read**
- **Microsoft.Compute/availabilitySets/write**
- **Microsoft.Compute/disks/beginGetAccess/action**
- **Microsoft.Compute/disks/delete**
- **Microsoft.Compute/disks/read**
- **Microsoft.Compute/disks/write**
- **Microsoft.Compute/galleries/images/read**
- **Microsoft.Compute/galleries/images/versions/read**
- **Microsoft.Compute/galleries/images/versions/write**
- **Microsoft.Compute/galleries/images/write**
- **Microsoft.Compute/galleries/read**
- **Microsoft.Compute/galleries/write**
- **Microsoft.Compute/snapshots/read**
- **Microsoft.Compute/snapshots/write**
- **Microsoft.Compute/snapshots/delete**
- **Microsoft.Compute/virtualMachines/delete**

- **Microsoft.Compute/virtualMachines/powerOff/action**
- **Microsoft.Compute/virtualMachines/read**
- **Microsoft.Compute/virtualMachines/write**

Example 7.3. Required permissions for creating identity management resources

- **Microsoft.ManagedIdentity/userAssignedIdentities/assign/action**
- **Microsoft.ManagedIdentity/userAssignedIdentities/read**
- **Microsoft.ManagedIdentity/userAssignedIdentities/write**

Example 7.4. Required permissions for creating network resources

- **Microsoft.Network/dnsZones/A/write**
- **Microsoft.Network/dnsZones/CNAME/write**
- **Microsoft.Network/dnszones/CNAME/read**
- **Microsoft.Network/dnszones/read**
- **Microsoft.Network/loadBalancers/backendAddressPools/join/action**
- **Microsoft.Network/loadBalancers/backendAddressPools/read**
- **Microsoft.Network/loadBalancers/backendAddressPools/write**
- **Microsoft.Network/loadBalancers/read**
- **Microsoft.Network/loadBalancers/write**
- **Microsoft.Network/networkInterfaces/delete**
- **Microsoft.Network/networkInterfaces/join/action**
- **Microsoft.Network/networkInterfaces/read**
- **Microsoft.Network/networkInterfaces/write**
- **Microsoft.Network/networkSecurityGroups/join/action**
- **Microsoft.Network/networkSecurityGroups/read**
- **Microsoft.Network/networkSecurityGroups/securityRules/delete**
- **Microsoft.Network/networkSecurityGroups/securityRules/read**
- **Microsoft.Network/networkSecurityGroups/securityRules/write**
- **Microsoft.Network/networkSecurityGroups/write**
- **Microsoft.Network/privateDnsZones/A/read**

- **Microsoft.Network/privateDnsZones/A/write**
- **Microsoft.Network/privateDnsZones/A/delete**
- **Microsoft.Network/privateDnsZones/SOA/read**
- **Microsoft.Network/privateDnsZones/read**
- **Microsoft.Network/privateDnsZones/virtualNetworkLinks/read**
- **Microsoft.Network/privateDnsZones/virtualNetworkLinks/write**
- **Microsoft.Network/privateDnsZones/write**
- **Microsoft.Network/publicIPAddresses/delete**
- **Microsoft.Network/publicIPAddresses/join/action**
- **Microsoft.Network/publicIPAddresses/read**
- **Microsoft.Network/publicIPAddresses/write**
- **Microsoft.Network/virtualNetworks/join/action**
- **Microsoft.Network/virtualNetworks/read**
- **Microsoft.Network/virtualNetworks/subnets/join/action**
- **Microsoft.Network/virtualNetworks/subnets/read**
- **Microsoft.Network/virtualNetworks/subnets/write**
- **Microsoft.Network/virtualNetworks/write**

NOTE

The following permissions are not required to create the private OpenShift Container Platform cluster on Azure.

- **Microsoft.Network/dnsZones/A/write**
- **Microsoft.Network/dnsZones/CNAME/write**
- **Microsoft.Network/dnszones/CNAME/read**
- **Microsoft.Network/dnszones/read**

Example 7.5. Required permissions for checking the health of resources

- **Microsoft.Resourcehealth/healthevent/Activated/action**
- **Microsoft.Resourcehealth/healthevent/InProgress/action**
- **Microsoft.Resourcehealth/healthevent/Pending/action**
- **Microsoft.Resourcehealth/healthevent/Resolved/action**

- **Microsoft.Resourcehealth/healthevent/Updated/action**

Example 7.6. Required permissions for creating a resource group

- **Microsoft.Resources/subscriptions/resourceGroups/read**
- **Microsoft.Resources/subscriptions/resourcegroups/write**

Example 7.7. Required permissions for creating resource tags

- **Microsoft.Resources/tags/write**

Example 7.8. Required permissions for creating storage resources

- **Microsoft.Storage/storageAccounts/blobServices/read**
- **Microsoft.Storage/storageAccounts/blobServices/containers/write**
- **Microsoft.Storage/storageAccounts/fileServices/read**
- **Microsoft.Storage/storageAccounts/fileServices/shares/read**
- **Microsoft.Storage/storageAccounts/fileServices/shares/write**
- **Microsoft.Storage/storageAccounts/fileServices/shares/delete**
- **Microsoft.Storage/storageAccounts/listKeys/action**
- **Microsoft.Storage/storageAccounts/read**
- **Microsoft.Storage/storageAccounts/write**

Example 7.9. Optional permissions for creating a private storage endpoint for the image registry

- **Microsoft.Network/privateEndpoints/write**
- **Microsoft.Network/privateEndpoints/read**
- **Microsoft.Network/privateEndpoints/privateDnsZoneGroups/write**
- **Microsoft.Network/privateEndpoints/privateDnsZoneGroups/read**
- **Microsoft.Network/privateDnsZones/join/action**
- **Microsoft.Storage/storageAccounts/PrivateEndpointConnectionsApproval/action**

Example 7.10. Optional permissions for creating marketplace virtual machine resources

- **Microsoft.MarketplaceOrdering/offertypes/publishers/offers/plans/agreements/read**

- **Microsoft.MarketplaceOrdering/offertypes/publishers/offers/plans/agreements/write**

Example 7.11. Optional permissions for creating compute resources

- **Microsoft.Compute/availabilitySets/delete**
- **Microsoft.Compute/images/read**
- **Microsoft.Compute/images/write**
- **Microsoft.Compute/images/delete**

Example 7.12. Optional permissions for enabling user-managed encryption

- **Microsoft.Compute/diskEncryptionSets/read**
- **Microsoft.Compute/diskEncryptionSets/write**
- **Microsoft.Compute/diskEncryptionSets/delete**
- **Microsoft.KeyVault/vaults/read**
- **Microsoft.KeyVault/vaults/write**
- **Microsoft.KeyVault/vaults/delete**
- **Microsoft.KeyVault/vaults/deploy/action**
- **Microsoft.KeyVault/vaults/keys/read**
- **Microsoft.KeyVault/vaults/keys/write**
- **Microsoft.Features/providers/features/register/action**

Example 7.13. Optional permissions for installing a cluster using the **NatGateway** outbound type

- **Microsoft.Network/natGateways/read**
- **Microsoft.Network/natGateways/write**

Example 7.14. Optional permissions for installing a private cluster with Azure Network Address Translation (NAT)

- **Microsoft.Network/natGateways/join/action**
- **Microsoft.Network/natGateways/read**
- **Microsoft.Network/natGateways/write**

Example 7.15. Optional permissions for installing a private cluster with Azure firewall

- **Microsoft.Network/azureFirewalls/applicationRuleCollections/write**
- **Microsoft.Network/azureFirewalls/read**
- **Microsoft.Network/azureFirewalls/write**
- **Microsoft.Network/routeTables/join/action**
- **Microsoft.Network/routeTables/read**
- **Microsoft.Network/routeTables/routes/read**
- **Microsoft.Network/routeTables/routes/write**
- **Microsoft.Network/routeTables/write**
- **Microsoft.Network/virtualNetworks/peer/action**
- **Microsoft.Network/virtualNetworks/virtualNetworkPeerings/read**
- **Microsoft.Network/virtualNetworks/virtualNetworkPeerings/write**

Example 7.16. Optional permission for running gather bootstrap

- **Microsoft.Compute/virtualMachines/instanceView/read**

The following permissions are required for deleting an OpenShift Container Platform cluster on Microsoft Azure. You can use the same permissions to delete a private OpenShift Container Platform cluster on Azure.

Example 7.17. Required permissions for deleting authorization resources

- **Microsoft.Authorization/roleAssignments/delete**

Example 7.18. Required permissions for deleting compute resources

- **Microsoft.Compute/disks/delete**
- **Microsoft.Compute/galleries/delete**
- **Microsoft.Compute/galleries/images/delete**
- **Microsoft.Compute/galleries/images/versions/delete**
- **Microsoft.Compute/virtualMachines/delete**

Example 7.19. Required permissions for deleting identity management resources

- **Microsoft.ManagedIdentity/userAssignedIdentities/delete**

Example 7.20. Required permissions for deleting network resources

- **Microsoft.Network/dnszones/read**
- **Microsoft.Network/dnsZones/A/read**
- **Microsoft.Network/dnsZones/A/delete**
- **Microsoft.Network/dnsZones/CNAME/read**
- **Microsoft.Network/dnsZones/CNAME/delete**
- **Microsoft.Network/loadBalancers/delete**
- **Microsoft.Network/networkInterfaces/delete**
- **Microsoft.Network/networkSecurityGroups/delete**
- **Microsoft.Network/privateDnsZones/read**
- **Microsoft.Network/privateDnsZones/A/read**
- **Microsoft.Network/privateDnsZones/delete**
- **Microsoft.Network/privateDnsZones/virtualNetworkLinks/delete**
- **Microsoft.Network/publicIPAddresses/delete**
- **Microsoft.Network/virtualNetworks/delete**

NOTE

The following permissions are not required to delete a private OpenShift Container Platform cluster on Azure.

- **Microsoft.Network/dnszones/read**
- **Microsoft.Network/dnsZones/A/read**
- **Microsoft.Network/dnsZones/A/delete**
- **Microsoft.Network/dnsZones/CNAME/read**
- **Microsoft.Network/dnsZones/CNAME/delete**

Example 7.21. Required permissions for checking the health of resources

- **Microsoft.Resourcehealth/healthevent/Activated/action**
- **Microsoft.Resourcehealth/healthevent/Resolved/action**
- **Microsoft.Resourcehealth/healthevent/Updated/action**

Example 7.22. Required permissions for deleting a resource group

- **Microsoft.Resources/subscriptions/resourcegroups/delete**

Example 7.23. Required permissions for deleting storage resources

- **Microsoft.Storage/storageAccounts/delete**
- **Microsoft.Storage/storageAccounts/listKeys/action**



NOTE

To install OpenShift Container Platform on Azure, you must scope the permissions to your subscription. Later, you can re-scope these permissions to the installer created resource group. If the public DNS zone is present in a different resource group, then the network DNS zone related permissions must always be applied to your subscription. By default, the OpenShift Container Platform installation program assigns the Azure identity the **Contributor** role.

You can scope all the permissions to your subscription when deleting an OpenShift Container Platform cluster.

7.2.5.3. Using Azure managed identities

The installation program requires an Azure identity to complete the installation. You can use either a system-assigned or user-assigned managed identity.

If you are unable to use a managed identity, you can use a service principal.

Procedure

1. If you are using a system-assigned managed identity, enable it on the virtual machine that you will run the installation program from.
2. If you are using a user-assigned managed identity:
 - a. Assign it to the virtual machine that you will run the installation program from.
 - b. Record its client ID. You require this value when installing the cluster.
For more information about viewing the details of a user-assigned managed identity, see the Microsoft Azure documentation for [listing user-assigned managed identities](#).
3. Verify that the required permissions are assigned to the managed identity.

7.2.5.4. Creating a service principal

The installation program requires an Azure identity to complete the installation. You can use a service principal.

If you are unable to use a service principal, you can use a managed identity.

Prerequisites

- You have installed or updated the [Azure CLI](#).

- You have an Azure subscription ID.
- If you are not going to assign the **Contributor** and **User Administrator Access** roles to the service principal, you have created a custom role with the required Azure permissions.

Procedure

1. Create the service principal for your account by running the following command:

```
$ az ad sp create-for-rbac --role <role_name> \ ❶
--name <service_principal> \ ❷
--scopes /subscriptions/<subscription_id> ❸
```

- ❶ Defines the role name. You can use the **Contributor** role, or you can specify a custom role which contains the necessary permissions.
- ❷ Defines the service principal name.
- ❸ Specifies the subscription ID.

Example output

```
Creating 'Contributor' role assignment under scope '/subscriptions/<subscription_id>'
The output includes credentials that you must protect. Be sure that you do not
include these credentials in your code or check the credentials into your source
control. For more information, see https://aka.ms/azadsp-cli
{
  "appId": "axxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
  "displayName": <service_principal>,
  "password": "00000000-0000-0000-0000-000000000000",
  "tenantId": "8xxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx"
}
```

2. Record the values of the **appId** and **password** parameters from the output. You require these values when installing the cluster.
3. If you applied the **Contributor** role to your service principal, assign the **User Administrator Access** role by running the following command:

```
$ az role assignment create --role "User Access Administrator" \
--assignee-object-id $(az ad sp show --id <appId> --query id -o tsv) ❶
--scope /subscriptions/<subscription_id> ❷
```

- ❶ Specify the **appId** parameter value for your service principal.
- ❷ Specifies the subscription ID.

Additional resources

- [About the Cloud Credential Operator](#)

7.2.6. Supported Azure Marketplace regions

Installing a cluster using the Azure Marketplace image is available to customers who purchase the offer in North America and EMEA.

While the offer must be purchased in North America or EMEA, you can deploy the cluster to any of the Azure public partitions that OpenShift Container Platform supports.



NOTE

Deploying a cluster using the Azure Marketplace image is not supported for the Azure Government regions.

7.2.7. Supported Azure regions

The installation program dynamically generates the list of available Microsoft Azure regions based on your subscription.

Supported Azure public regions

- **australiacentral** (Australia Central)
- **australiaeast** (Australia East)
- **australiasoutheast** (Australia South East)
- **brazilsouth** (Brazil South)
- **canadacentral** (Canada Central)
- **canadaeast** (Canada East)
- **centralindia** (Central India)
- **centralus** (Central US)
- **eastasia** (East Asia)
- **eastus** (East US)
- **eastus2** (East US 2)
- **francecentral** (France Central)
- **germanywestcentral** (Germany West Central)
- **israelcentral** (Israel Central)
- **italynorth** (Italy North)
- **japaneast** (Japan East)
- **japanwest** (Japan West)
- **koreacentral** (Korea Central)
- **koreasouth** (Korea South)
- **mexicocentral** (Mexico Central)

- **northcentralus** (North Central US)
- **northeurope** (North Europe)
- **norwayeast** (Norway East)
- **polandcentral** (Poland Central)
- **qatarcentral** (Qatar Central)
- **southafricanorth** (South Africa North)
- **southcentralus** (South Central US)
- **southeastasia** (Southeast Asia)
- **southindia** (South India)
- **swedencentral** (Sweden Central)
- **switzerlandnorth** (Switzerland North)
- **uaenorth** (UAE North)
- **uksouth** (UK South)
- **ukwest** (UK West)
- **westcentralus** (West Central US)
- **westeurope** (West Europe)
- **westindia** (West India)
- **westus** (West US)
- **westus2** (West US 2)
- **westus3** (West US 3)

Supported Azure Government regions

Support for the following Microsoft Azure Government (MAG) regions was added in OpenShift Container Platform version 4.6:

- **usgovtexas** (US Gov Texas)
- **usgovvirginia** (US Gov Virginia)

You can reference all available MAG regions in the [Azure documentation](#). Other provided MAG regions are expected to work with OpenShift Container Platform, but have not been tested.

7.2.8. Next steps

- Install an OpenShift Container Platform cluster on Azure. You can [install a customized cluster](#) or [quickly install a cluster](#) with default options.

7.3. ENABLING USER-MANAGED ENCRYPTION FOR AZURE

In OpenShift Container Platform version 4.16, you can install a cluster with a user-managed encryption key in Azure. To enable this feature, you can prepare an Azure DiskEncryptionSet before installation, modify the `install-config.yaml` file, and then complete the installation.

7.3.1. Preparing an Azure Disk Encryption Set

The OpenShift Container Platform installer can use an existing Disk Encryption Set with a user-managed key. To enable this feature, you can create a Disk Encryption Set in Azure and provide the key to the installer.

Procedure

1. Set the following environment variables for the Azure resource group by running the following command:

```
$ export RESOURCEGROUP="<resource_group>" 1
LOCATION="<location>" 2
```

- 1** Specifies the name of the Azure resource group where you will create the Disk Encryption Set and encryption key. To avoid losing access to your keys after destroying the cluster, you should create the Disk Encryption Set in a different resource group than the resource group where you install the cluster.

- 2** Specifies the Azure location where you will create the resource group.

2. Set the following environment variables for the Azure Key Vault and Disk Encryption Set by running the following command:

```
$ export KEYVAULT_NAME="<keyvault_name>" 1
KEYVAULT_KEY_NAME="<keyvault_key_name>" 2
DISK_ENCRYPTION_SET_NAME="<disk_encryption_set_name>" 3
```

- 1** Specifies the name of the Azure Key Vault you will create.

- 2** Specifies the name of the encryption key you will create.

- 3** Specifies the name of the disk encryption set you will create.

3. Set the environment variable for the ID of your Azure Service Principal by running the following command:

```
$ export CLUSTER_SP_ID="<service_principal_id>" 1
```

- 1** Specifies the ID of the service principal you will use for this installation.

4. Enable host-level encryption in Azure by running the following commands:

```
$ az feature register --namespace "Microsoft.Compute" --name "EncryptionAtHost"
```

```
$ az feature show --namespace Microsoft.Compute --name EncryptionAtHost
```

-


```
$ az provider register -n Microsoft.Compute
```

5. Create an Azure Resource Group to hold the disk encryption set and associated resources by running the following command:

```
$ az group create --name $RESOURCEGROUP --location $LOCATION
```

6. Create an Azure key vault by running the following command:

```
$ az keyvault create -n $KEYVAULT_NAME -g $RESOURCEGROUP -l $LOCATION \
  --enable-purge-protection true
```

7. Create an encryption key in the key vault by running the following command:

```
$ az keyvault key create --vault-name $KEYVAULT_NAME -n $KEYVAULT_KEY_NAME \
  --protection software
```

8. Capture the ID of the key vault by running the following command:

```
$ KEYVAULT_ID=$(az keyvault show --name $KEYVAULT_NAME --query "[id]" -o tsv)
```

9. Capture the key URL in the key vault by running the following command:

```
$ KEYVAULT_KEY_URL=$(az keyvault key show --vault-name $KEYVAULT_NAME --name \
  $KEYVAULT_KEY_NAME --query "[key.kid]" -o tsv)
```

10. Create a disk encryption set by running the following command:

```
$ az disk-encryption-set create -n $DISK_ENCRYPTION_SET_NAME -l $LOCATION -g \
  $RESOURCEGROUP --source-vault $KEYVAULT_ID --key-url $KEYVAULT_KEY_URL
```

11. Grant the DiskEncryptionSet resource access to the key vault by running the following commands:

```
$ DES_IDENTITY=$(az disk-encryption-set show -n $DISK_ENCRYPTION_SET_NAME -g \
  $RESOURCEGROUP --query "[identity.principalId]" -o tsv)
```

```
$ az keyvault set-policy -n $KEYVAULT_NAME -g $RESOURCEGROUP --object-id \
  $DES_IDENTITY --key-permissions wrapkey unwrapkey get
```

12. Grant the Azure Service Principal permission to read the DiskEncryptionSet by running the following commands:

```
$ DES_RESOURCE_ID=$(az disk-encryption-set show -n \
  $DISK_ENCRYPTION_SET_NAME -g \
  $RESOURCEGROUP --query "[id]" -o tsv)
```

```
$ az role assignment create --assignee $CLUSTER_SP_ID --role "<reader_role>" \
  --scope $DES_RESOURCE_ID -o jsonc 1
```

- 1 Specifies an Azure role with read permissions to the disk encryption set. You can use the **Owner** role or a custom role with the necessary permissions.

7.3.2. Next steps

- Install an OpenShift Container Platform cluster:
 - [Install a cluster with customizations on installer-provisioned infrastructure](#)
 - [Install a cluster with network customizations on installer-provisioned infrastructure](#)
 - [Install a cluster into an existing VNet on installer-provisioned infrastructure](#)
 - [Install a private cluster on installer-provisioned infrastructure](#)
 - [Install a cluster into an government region on installer-provisioned infrastructure](#)

7.4. INSTALLING A CLUSTER QUICKLY ON AZURE

In OpenShift Container Platform version 4.16, you can install a cluster on Microsoft Azure that uses the default configuration options.

7.4.1. Prerequisites


- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [configured an Azure account](#) to host the cluster and determined the tested and validated region to deploy the cluster to.
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.

7.4.2. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.



IMPORTANT


If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

7.4.3. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.


After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.



NOTE

You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.



NOTE

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.


- View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

- Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.



NOTE

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- If the `ssh-agent` process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```



NOTE

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

- Add your SSH private key to the `ssh-agent`:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

7.4.4. Obtaining the installation program


Before you install OpenShift Container Platform, download the installation file on the host you are using for installation.

Prerequisites

- You have a computer that runs Linux or macOS, with 500 MB of local disk space.


Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select your infrastructure provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.



IMPORTANT

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.



IMPORTANT

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.


4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#). This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

7.4.5. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have the OpenShift Container Platform installation program and the pull secret for your cluster.

- You have an Azure subscription ID and tenant ID.
- You have the application ID and password of a service principal.

Procedure

1. Optional: If you have run the installation program on this computer before, and want to use an alternative service principal, go to the `~/.azure/` directory and delete the **osServicePrincipal.json** configuration file.
Deleting this file prevents the installation program from automatically reusing subscription and authentication values from a previous installation.
2. Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

- 1 For **<installation_directory>**, specify the directory name to store the files that the installation program creates.
- 2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

When specifying the directory:

- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
- Use an empty directory. Some installation assets, such as bootstrap X.509 certificates, have short expiration intervals, therefore you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

3. Provide values at the prompts:
 - a. Optional: Select an SSH key to use to access your cluster machines.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- b. Select **azure** as the platform to target.
If the installation program cannot locate the **osServicePrincipal.json** configuration file from a previous installation, you are prompted for Azure subscription and authentication values.
- c. Specify the following Azure parameter values for your subscription and service principal:
 - **azure subscription id** Enter the subscription ID to use for the cluster.
 - **azure tenant id** Enter the tenant ID.

- **azure service principal client id** Enter its application ID.
 - **azure service principal client secret** Enter its password.
- d. Select the region to deploy the cluster to.
- e. Select the base domain to deploy the cluster to. The base domain corresponds to the Azure DNS Zone that you created for your cluster.
- f. Enter a descriptive name for your cluster.



IMPORTANT

All Azure resources that are available through public endpoints are subject to resource name restrictions, and you cannot create resources that use certain terms. For a list of terms that Azure restricts, see [Resolve reserved resource name errors](#) in the Azure documentation.


- g. Paste the [pull secret from Red Hat OpenShift Cluster Manager](#) .

If previously not detected, the installation program creates an **osServicePrincipal.json** configuration file and stores this file in the `~/.azure/` directory on your computer. This ensures that the installation program can load the profile when it is creating an OpenShift Container Platform cluster on the target platform.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to `<installation_directory>/openshift_install.log`.



IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

7.4.6. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.



IMPORTANT

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**. To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```


-

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:


```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.



NOTE

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your PATH.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

7.4.7. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

Additional resources

- See [Accessing the web console](#) for more details about accessing and understanding the OpenShift Container Platform web console.

7.4.8. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service

7.4.9. Next steps

- [Customize your cluster.](#)
- If necessary, you can [opt out of remote health reporting](#) .

7.5. INSTALLING A CLUSTER ON AZURE WITH CUSTOMIZATIONS

In OpenShift Container Platform version 4.16, you can install a customized cluster on infrastructure that the installation program provisions on Microsoft Azure. To customize the installation, you modify parameters in the **install-config.yaml** file before you install the cluster.

7.5.1. Prerequisites


- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [configured an Azure account](#) to host the cluster and determined the tested and validated region to deploy the cluster to.
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.
- If you use customer-managed encryption keys, you [prepared your Azure environment for encryption](#).

7.5.2. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.



IMPORTANT

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

7.5.3. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes

through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.



NOTE

You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.



NOTE

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.

**NOTE**


On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```

**NOTE**

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

7.5.4. Using the Azure Marketplace offering

Using the Azure Marketplace offering lets you deploy an OpenShift Container Platform cluster, which is billed on pay-per-use basis (hourly, per core) through Azure, while still being supported directly by Red Hat.

To deploy an OpenShift Container Platform cluster using the Azure Marketplace offering, you must first obtain the Azure Marketplace image. The installation program uses this image to deploy worker or control plane nodes. When obtaining your image, consider the following:

- While the images are the same, the Azure Marketplace publisher is different depending on your region. If you are located in North America, specify **redhat** as the publisher. If you are located in EMEA, specify **redhat-limited** as the publisher.
- The offer includes a **rh-ocp-worker** SKU and a **rh-ocp-worker-gen1** SKU. The **rh-ocp-worker** SKU represents a Hyper-V generation version 2 VM image. The default instance types used in OpenShift Container Platform are version 2 compatible. If you plan to use an instance type that

is only version 1 compatible, use the image associated with the **rh-ocp-worker-gen1** SKU. The **rh-ocp-worker-gen1** SKU represents a Hyper-V version 1 VM image.



IMPORTANT

Installing images with the Azure marketplace is not supported on clusters with 64-bit ARM instances.

Prerequisites

- You have installed the Azure CLI client (**az**).
- Your Azure account is entitled for the offer and you have logged into this account with the Azure CLI client.

Procedure

1. Display all of the available OpenShift Container Platform images by running one of the following commands:

- North America:

```
$ az vm image list --all --offer rh-ocp-worker --publisher redhat -o table
```

Example output

Offer	Publisher	Sku	Urn	Version
rh-ocp-worker	RedHat	rh-ocp-worker	RedHat:rh-ocp-worker:rh-ocp-worker:413.92.2023101700	413.92.2023101700
rh-ocp-worker-gen1	RedHat	rh-ocp-worker-gen1	RedHat:rh-ocp-worker:rh-ocp-worker-gen1:413.92.2023101700	413.92.2023101700

- EMEA:

```
$ az vm image list --all --offer rh-ocp-worker --publisher redhat-limited -o table
```

Example output

Offer	Publisher	Sku	Urn	Version
rh-ocp-worker	redhat-limited	rh-ocp-worker	redhat-limited:rh-ocp-worker:rh-ocp-worker:413.92.2023101700	413.92.2023101700
rh-ocp-worker-gen1	redhat-limited	rh-ocp-worker-gen1	redhat-limited:rh-ocp-worker:rh-ocp-worker-gen1:413.92.2023101700	413.92.2023101700



NOTE

Use the latest image that is available for compute and control plane nodes. If required, your VMs are automatically upgraded as part of the installation process.

2. Inspect the image for your offer by running one of the following commands:

- North America:

```
$ az vm image show --urn redhat:rh-ocp-worker:rh-ocp-worker:<version>
```

- EMEA:

```
$ az vm image show --urn redhat-limited:rh-ocp-worker:rh-ocp-worker:<version>
```

3. Review the terms of the offer by running one of the following commands:

- North America:

```
$ az vm image terms show --urn redhat:rh-ocp-worker:rh-ocp-worker:<version>
```

- EMEA:

```
$ az vm image terms show --urn redhat-limited:rh-ocp-worker:rh-ocp-worker:<version>
```

4. Accept the terms of the offering by running one of the following commands:

- North America:

```
$ az vm image terms accept --urn redhat:rh-ocp-worker:rh-ocp-worker:<version>
```

- EMEA:

```
$ az vm image terms accept --urn redhat-limited:rh-ocp-worker:rh-ocp-worker:<version>
```

5. Record the image details of your offer. You must update the **compute** section in the **install-config.yaml** file with values for **publisher**, **offer**, **sku**, and **version** before deploying the cluster. You may also update the **controlPlane** section to deploy control plane machines with the specified image details, or the **defaultMachinePlatform** section to deploy both control plane and compute machines with the specified image details. Use the latest available image for control plane and compute nodes.

Sample **install-config.yaml** file with the Azure Marketplace compute nodes

```
apiVersion: v1
baseDomain: example.com
compute:
- hyperthreading: Enabled
  name: worker
  platform:
    azure:
      type: Standard_D4s_v5
      osImage:
        publisher: redhat
        offer: rh-ocp-worker
        sku: rh-ocp-worker
        version: 413.92.2023101700
  replicas: 3
```

7.5.5. Obtaining the installation program

Before you install OpenShift Container Platform, download the installation file on the host you are using for installation.

Prerequisites

- You have a computer that runs Linux or macOS, with 500 MB of local disk space.

Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select your infrastructure provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.



IMPORTANT

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.



IMPORTANT

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.

4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#) . This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

7.5.6. Creating the installation configuration file

You can customize the OpenShift Container Platform cluster you install on Microsoft Azure.

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have an Azure subscription ID and tenant ID.

- If you are installing the cluster using a service principal, you have its application ID and password.
- If you are installing the cluster using a system-assigned managed identity, you have enabled it on the virtual machine that you will run the installation program from.
- If you are installing the cluster using a user-assigned managed identity, you have met these prerequisites:
 - You have its client ID.
 - You have assigned it to the virtual machine that you will run the installation program from.

Procedure

1. Optional: If you have run the installation program on this computer before, and want to use an alternative service principal or managed identity, go to the `~/.azure/` directory and delete the **osServicePrincipal.json** configuration file.
Deleting this file prevents the installation program from automatically reusing subscription and authentication values from a previous installation.

2. Create the **install-config.yaml** file.

- a. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install create install-config --dir <installation_directory> 1
```

- 1 For **<installation_directory>**, specify the directory name to store the files that the installation program creates.

When specifying the directory:

- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
- Use an empty directory. Some installation assets, such as bootstrap X.509 certificates, have short expiration intervals, therefore you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

- b. At the prompts, provide the configuration details for your cloud:

- i. Optional: Select an SSH key to use to access your cluster machines.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- ii. Select **azure** as the platform to target.

If the installation program cannot locate the **osServicePrincipal.json** configuration file from a previous installation, you are prompted for Azure subscription and authentication values.

- iii. Enter the following Azure parameter values for your subscription:
 - **azure subscription id** Enter the subscription ID to use for the cluster.
 - **azure tenant id** Enter the tenant ID.
- iv. Depending on the Azure identity you are using to deploy the cluster, do one of the following when prompted for the **azure service principal client id**
 - If you are using a service principal, enter its application ID.
 - If you are using a system-assigned managed identity, leave this value blank.
 - If you are using a user-assigned managed identity, specify its client ID.
- v. Depending on the Azure identity you are using to deploy the cluster, do one of the following when prompted for the **azure service principal client secret**
 - If you are using a service principal, enter its password.
 - If you are using a system-assigned managed identity, leave this value blank.
 - If you are using a user-assigned managed identity, leave this value blank.
- vi. Select the region to deploy the cluster to.
- vii. Select the base domain to deploy the cluster to. The base domain corresponds to the Azure DNS Zone that you created for your cluster.
- viii. Enter a descriptive name for your cluster.



IMPORTANT

All Azure resources that are available through public endpoints are subject to resource name restrictions, and you cannot create resources that use certain terms. For a list of terms that Azure restricts, see [Resolve reserved resource name errors](#) in the Azure documentation.


3. Modify the **install-config.yaml** file. You can find more information about the available parameters in the "Installation configuration parameters" section.



NOTE

If you are installing a three-node cluster, be sure to set the **compute.replicas** parameter to **0**. This ensures that the cluster's control planes are schedulable. For more information, see "Installing a three-node cluster on Azure".

4. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

If previously not detected, the installation program creates an **osServicePrincipal.json** configuration file and stores this file in the `~/.azure/` directory on your computer. This ensures that the installation program can load the profile when it is creating an OpenShift Container Platform cluster on the target platform.

Additional resources

- [Installation configuration parameters for Azure](#)

7.5.6.1. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 7.1. Minimum resource requirements

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Bootstrap	RHCOS	4	16 GB	100 GB	300
Control plane	RHCOS	4	16 GB	100 GB	300
Compute	RHCOS, RHEL 8.6 and later [3]	2	8 GB	100 GB	300

1. One vCPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio: $(\text{threads per core} \times \text{cores}) \times \text{sockets} = \text{vCPUs}$.
2. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes which require a 10 ms p99 fsync duration. Note that on many cloud platforms, storage size and IOPS scale together, so you might need to over-allocate storage volume to obtain sufficient performance.
3. As with all user-provisioned installations, if you choose to use RHEL compute machines in your cluster, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Use of RHEL 7 compute machines is deprecated and has been removed in OpenShift Container Platform 4.10 and later.



NOTE

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).



IMPORTANT

You are required to use Azure virtual machines that have the **premiumIO** parameter set to **true**.

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

7.5.6.2. Tested instance types for Azure

The following Microsoft Azure instance types have been tested with OpenShift Container Platform.

Example 7.24. Machine types based on 64-bit x86 architecture

- **c4.***
- **c5.***
- **c5a.***
- **i3.***
- **m4.***
- **m5.***
- **m5a.***
- **m6a.***
- **m6i.***
- **r4.***
- **r5.***

- **r5a.***
- **r6i.***
- **t3.***
- **t3a.***

7.5.6.3. Tested instance types for Azure on 64-bit ARM infrastructures

The following Microsoft Azure ARM64 instance types have been tested with OpenShift Container Platform.

Example 7.25. Machine types based on 64-bit ARM architecture

- **c6g.***
- **m6g.***

7.5.6.4. Enabling trusted launch for Azure VMs

You can enable two trusted launch features when installing your cluster on Azure: [secure boot](#) and [virtualized Trusted Platform Modules](#).

See the Azure documentation about [virtual machine sizes](#) to learn what sizes of virtual machines support these features.



IMPORTANT

Trusted launch is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

Prerequisites

- You have created an **install-config.yaml** file.

Procedure

- Use a text editor to edit the **install-config.yaml** file prior to deploying your cluster and add the following stanza:

```
controlPlane: ❶
platform:
  azure:
    settings:
      securityType: TrustedLaunch ❷
```

```


trustedLaunch:
  uefiSettings:
    secureBoot: Enabled 3
    virtualizedTrustedPlatformModule: Enabled 4

```

- 1** Specify **controlPlane.platform.azure** or **compute.platform.azure** to enable trusted launch on only control plane or compute nodes respectively. Specify **platform.azure.defaultMachinePlatform** to enable trusted launch on all nodes.
- 2** Enable trusted launch features.
- 3** Enable secure boot. For more information, see the Azure documentation about [secure boot](#).
- 4** Enable the virtualized Trusted Platform Module. For more information, see the Azure documentation about [virtualized Trusted Platform Modules](#).

7.5.6.5. Enabling confidential VMs

You can enable confidential VMs when installing your cluster. You can enable confidential VMs for compute nodes, control plane nodes, or all nodes.




IMPORTANT

Using confidential VMs is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

You can use confidential VMs with the following VM sizes:

- DCasv5-series
- DCadsv5-series
- ECasv5-series
- ECadsv5-series



IMPORTANT

Confidential VMs are currently not supported on 64-bit ARM architectures.

Prerequisites

- You have created an **install-config.yaml** file.

Procedure

- Use a text editor to edit the **install-config.yaml** file prior to deploying your cluster and add the following stanza:

```
controlPlane: ❶
  platform:
    azure:
      settings:
        securityType: ConfidentialVM ❷
        confidentialVM:
          uefiSettings:
            secureBoot: Enabled ❸
            virtualizedTrustedPlatformModule: Enabled ❹
      osDisk:
        securityProfile:
          securityEncryptionType: VMGuestStateOnly ❺
```

- ❶ Specify **controlPlane.platform.azure** or **compute.platform.azure** to deploy confidential VMs on only control plane or compute nodes respectively. Specify **platform.azure.defaultMachinePlatform** to deploy confidential VMs on all nodes.
- ❷ Enable confidential VMs.
- ❸ Enable secure boot. For more information, see the Azure documentation about [secure boot](#).
- ❹ Enable the virtualized Trusted Platform Module. For more information, see the Azure documentation about [virtualized Trusted Platform Modules](#).
- ❺ Specify **VMGuestStateOnly** to encrypt the VM guest state.

7.5.6.6. Sample customized install-config.yaml file for Azure

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.



IMPORTANT

This sample YAML file is provided for reference only. You must obtain your **install-config.yaml** file by using the installation program and modify it.

```
apiVersion: v1
baseDomain: example.com ❶
controlPlane: ❷
  hyperthreading: Enabled ❸ ❹
  name: master
  platform:
    azure:
      encryptionAtHost: true
      ultraSSDCapability: Enabled
      osDisk:
        diskSizeGB: 1024 ❺
        diskType: Premium_LRS
        diskEncryptionSet:
```

```
    resourceGroup: disk_encryption_set_resource_group
    name: disk_encryption_set_name
    subscriptionId: secondary_subscription_id
  osImage:
    publisher: example_publisher_name
    offer: example_image_offer
    sku: example_offer_sku
    version: example_image_version
    type: Standard_D8s_v3
  replicas: 3
compute: 6
- hyperthreading: Enabled 7
  name: worker
  platform:
    azure:
      ultraSSDCapability: Enabled
      type: Standard_D2s_v3
      encryptionAtHost: true
    osDisk:
      diskSizeGB: 512 8
      diskType: Standard_LRS
      diskEncryptionSet:
        resourceGroup: disk_encryption_set_resource_group
        name: disk_encryption_set_name
        subscriptionId: secondary_subscription_id
    osImage:
      publisher: example_publisher_name
      offer: example_image_offer
      sku: example_offer_sku
      version: example_image_version
  zones: 9
  - "1"
  - "2"
  - "3"
  replicas: 5
metadata:
  name: test-cluster 10
networking:
  clusterNetwork:
    - cidr: 10.128.0.0/14
      hostPrefix: 23
  machineNetwork:
    - cidr: 10.0.0.0/16
  networkType: OVNKubernetes 11
  serviceNetwork:
    - 172.30.0.0/16
platform:
  azure:
    defaultMachinePlatform:
      osImage: 12
      publisher: example_publisher_name
      offer: example_image_offer
      sku: example_offer_sku
      version: example_image_version
      ultraSSDCapability: Enabled
```



```

baseDomainResourceGroupName: resource_group 13
region: centralus 14
resourceGroupName: existing_resource_group 15
outboundType: Loadbalancer
cloudName: AzurePublicCloud
pullSecret: '{"auths": ...}' 16
fips: false 17
sshKey: ssh-ed25519 AAAA... 18

```

- 1 10 14 16 Required. The installation program prompts you for this value.
- 2 6 If you do not provide these parameters and values, the installation program provides the default value.
- 3 7 The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.
- 4 Whether to enable or disable simultaneous multithreading, or **hypertreading**. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores. You can disable it by setting the parameter value to **Disabled**. If you disable simultaneous multithreading in some cluster machines, you must disable it in all cluster machines.



IMPORTANT

If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance. Use larger virtual machine types, such as **Standard_D8s_v3**, for your machines if you disable simultaneous multithreading.

- 5 8 You can specify the size of the disk to use in GB. Minimum recommendation for control plane nodes is 1024 GB.
- 9 Specify a list of zones to deploy your machines to. For high availability, specify at least two zones.
- 11 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 12 Optional: A custom Red Hat Enterprise Linux CoreOS (RHCOS) image that should be used to boot control plane and compute machines. The **publisher**, **offer**, **sku**, and **version** parameters under **platform.azure.defaultMachinePlatform.osImage** apply to both control plane and compute machines. If the parameters under **controlPlane.platform.azure.osImage** or **compute.platform.azure.osImage** are set, they override the **platform.azure.defaultMachinePlatform.osImage** parameters.
- 13 Specify the name of the resource group that contains the DNS zone for your base domain.
- 15 Specify the name of an already existing resource group to install your cluster to. If undefined, a new resource group is created for the cluster.
- 17 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container




IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 18** You can optionally provide the **sshKey** value that you use to access the machines in your cluster.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

7.5.6.7. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

- Edit your **install-config.yaml** file and add the proxy settings. For example:


```
apiVersion: v1
baseDomain: my.domain.com
```

```

proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5


```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.



NOTE

The installation program does not support the proxy **readinessEndpoints** field.



NOTE

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.

**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

Additional resources

- For more details about Accelerated Networking, see [Accelerated Networking for Microsoft Azure VMs](#).

7.5.7. Configuring user-defined tags for Azure

In OpenShift Container Platform, you can use the tags for grouping resources and for managing resource access and cost. You can define the tags on the Azure resources in the **install-config.yaml** file only during OpenShift Container Platform cluster creation. You cannot modify the user-defined tags after cluster creation.

Support for user-defined tags is available only for the resources created in the Azure Public Cloud. User-defined tags are not supported for the OpenShift Container Platform clusters upgraded to OpenShift Container Platform 4.16.

User-defined and OpenShift Container Platform specific tags are applied only to the resources created by the OpenShift Container Platform installer and its core operators such as Machine api provider azure Operator, Cluster Ingress Operator, Cluster Image Registry Operator.

By default, OpenShift Container Platform installer attaches the OpenShift Container Platform tags to the Azure resources. These OpenShift Container Platform tags are not accessible for the users.

You can use the **.platform.azure.userTags** field in the **install-config.yaml** file to define the list of user-defined tags as shown in the following **install-config.yaml** file.

Sample install-config.yaml file

```

additionalTrustBundlePolicy: Proxyonly 1
apiVersion: v1
baseDomain: catchall.azure.devcluster.openshift.com 2
compute: 3
- architecture: amd64
  hyperthreading: Enabled 4
  name: worker
  platform: {}
  replicas: 3
controlPlane: 5
  architecture: amd64
  hyperthreading: Enabled 6
  name: master
  platform: {}
  replicas: 3
metadata:
  creationTimestamp: null
  name: user 7
networking:
  clusterNetwork:
  - cidr: 10.128.0.0/14
    hostPrefix: 23

```

```

machineNetwork:
- cidr: 10.0.0.0/16
networkType: OVNKubernetes 8
serviceNetwork:
- 172.30.0.0/16
platform:
azure:
baseDomainResourceGroupName: os4-common 9
cloudName: AzurePublicCloud 10
outboundType: Loadbalancer
region: southindia 11
userTags: 12
createdBy: user
environment: dev

```

- 1 Defines the trust bundle policy.
- 2 Required. The **baseDomain** parameter specifies the base domain of your cloud provider. The installation program prompts you for this value.
- 3 The configuration for the machines that comprise compute. The **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -. If you do not provide these parameters and values, the installation program provides the default value.
- 4 To enable or disable simultaneous multithreading, or **hyperthreading**. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores. You can disable it by setting the parameter value to **Disabled**. If you disable simultaneous multithreading in some cluster machines, you must disable it in all cluster machines.
- 5 The configuration for the machines that comprise the control plane. The **controlPlane** section is a single mapping. The first line of the **controlPlane** section must not begin with a hyphen, -. You can use only one control plane pool. If you do not provide these parameters and values, the installation program provides the default value.
- 6 To enable or disable simultaneous multithreading, or **hyperthreading**. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores. You can disable it by setting the parameter value to **Disabled**. If you disable simultaneous multithreading in some cluster machines, you must disable it in all cluster machines.
- 7 The installation program prompts you for this value.
- 8 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 9 Specifies the resource group for the base domain of the Azure DNS zone.
- 10 Specifies the name of the Azure cloud environment. You can use the **cloudName** field to configure the Azure SDK with the Azure API endpoints. If you do not provide value, the default value is Azure Public Cloud.
- 11 Required. Specifies the name of the Azure region that hosts your cluster. The installation program prompts you for this value.
- 12 Defines the additional keys and values that the installation program adds as tags to all Azure resources that it creates.

The user-defined tags have the following limitations:

- A tag key can have a maximum of 128 characters.
- A tag key must begin with a letter, end with a letter, number or underscore, and can contain only letters, numbers, underscores, periods, and hyphens.
- Tag keys are case-insensitive.
- Tag keys cannot be **name**. It cannot have prefixes such as **kubernetes.io**, **openshift.io**, **microsoft**, **azure**, and **windows**.
- A tag value can have a maximum of 256 characters.
- You can configure a maximum of 10 tags for resource group and resources.

For more information about Azure tags, see [Azure user-defined tags](#)

7.5.8. Querying user-defined tags for Azure

After creating the OpenShift Container Platform cluster, you can access the list of defined tags for the Azure resources. The format of the OpenShift Container Platform tags is **kubernetes.io_cluster.<cluster_id>:owned**. The **cluster_id** parameter is the value of **.status.infrastructureName** present in **config.openshift.io/Infrastructure**.

- Query the tags defined for Azure resources by running the following command:

```
$ oc get infrastructures.config.openshift.io cluster -o=jsonpath-as-
json='{.status.platformStatus.azure.resourceTags}'
```

Example output

```
[
  [
    {
      "key": "createdBy",
      "value": "user"
    },
    {
      "key": "environment",
      "value": "dev"
    }
  ]
]
```

7.5.9. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.



IMPORTANT

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**.
To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:


```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.



NOTE

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your PATH.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

7.5.10. Alternatives to storing administrator-level secrets in the kube-system project

By default, administrator secrets are stored in the **kube-system** project. If you configured the **credentialsMode** parameter in the **install-config.yaml** file to **Manual**, you must use one of the following alternatives:

- To manage long-term cloud credentials manually, follow the procedure in [Manually creating long-term credentials](#).
- To implement short-term credentials that are managed outside the cluster for individual components, follow the procedures in [Configuring an Azure cluster to use short-term credentials](#).

7.5.10.1. Manually creating long-term credentials

The Cloud Credential Operator (CCO) can be put into manual mode prior to installation in environments where the cloud identity and access management (IAM) APIs are not reachable, or the administrator prefers not to store an administrator-level credential secret in the cluster **kube-system** namespace.

Procedure

1. If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```
apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...
```

2. If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

3. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

4. Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included \1
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \2
  --to=<path_to_directory_for_credentials_requests> \3
```

- 1 The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2 Specify the location of the **install-config.yaml** file.
- 3 Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

This command creates a YAML file for each **CredentialsRequest** object.

Sample CredentialsRequest object

```
apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
```

```

kind: AzureProviderSpec
roleBindings:
- role: Contributor
...

```

5. Create YAML files for secrets in the **openshift-install** manifests directory that you generated previously. The secrets must be stored using the namespace and secret name defined in the **spec.secretRef** for each **CredentialsRequest** object.

Sample CredentialsRequest object with secrets

```

apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: AzureProviderSpec
    roleBindings:
      - role: Contributor
    ...
  secretRef:
    name: <component_secret>
    namespace: <component_namespace>
...


```

Sample Secret object

```

apiVersion: v1
kind: Secret
metadata:
  name: <component_secret>
  namespace: <component_namespace>
data:
  azure_subscription_id: <base64_encoded_azure_subscription_id>
  azure_client_id: <base64_encoded_azure_client_id>
  azure_client_secret: <base64_encoded_azure_client_secret>
  azure_tenant_id: <base64_encoded_azure_tenant_id>
  azure_resource_prefix: <base64_encoded_azure_resource_prefix>
  azure_resourcegroup: <base64_encoded_azure_resourcegroup>
  azure_region: <base64_encoded_azure_region>

```



IMPORTANT

Before upgrading a cluster that uses manually maintained credentials, you must ensure that the CCO is in an upgradeable state.

7.5.10.2. Configuring an Azure cluster to use short-term credentials

To install a cluster that uses Microsoft Entra Workload ID, you must configure the Cloud Credential Operator utility and create the required Azure resources for your cluster.

7.5.10.2.1. Configuring the Cloud Credential Operator utility

To create and manage cloud credentials from outside of the cluster when the Cloud Credential Operator (CCO) is operating in manual mode, extract and prepare the CCO utility (**ccoctl**) binary.



NOTE

The **ccoctl** utility is a Linux binary that must run in a Linux environment.

Prerequisites

- You have access to an OpenShift Container Platform account with cluster administrator access.
- You have installed the OpenShift CLI (**oc**).
- You have created a global Microsoft Azure account for the **ccoctl** utility to use with the following permissions:

Example 7.26. Required Azure permissions

- Microsoft.Resources/subscriptions/resourceGroups/read
- Microsoft.Resources/subscriptions/resourceGroups/write
- Microsoft.Resources/subscriptions/resourceGroups/delete
- Microsoft.Authorization/roleAssignments/read
- Microsoft.Authorization/roleAssignments/delete
- Microsoft.Authorization/roleAssignments/write
- Microsoft.Authorization/roleDefinitions/read
- Microsoft.Authorization/roleDefinitions/write
- Microsoft.Authorization/roleDefinitions/delete
- Microsoft.Storage/storageAccounts/listkeys/action
- Microsoft.Storage/storageAccounts/delete
- Microsoft.Storage/storageAccounts/read
- Microsoft.Storage/storageAccounts/write
- Microsoft.Storage/storageAccounts/blobServices/containers/write
- Microsoft.Storage/storageAccounts/blobServices/containers/delete
- Microsoft.Storage/storageAccounts/blobServices/containers/read
- Microsoft.ManagedIdentity/userAssignedIdentities/delete
- Microsoft.ManagedIdentity/userAssignedIdentities/read
- Microsoft.ManagedIdentity/userAssignedIdentities/write

- Microsoft.ManagedIdentity/userAssignedIdentities/federatedIdentityCredentials/read
- Microsoft.ManagedIdentity/userAssignedIdentities/federatedIdentityCredentials/write
- Microsoft.ManagedIdentity/userAssignedIdentities/federatedIdentityCredentials/delete
- Microsoft.Storage/register/action
- Microsoft.ManagedIdentity/register/action


Procedure

1. Set a variable for the OpenShift Container Platform release image by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

2. Obtain the CCO container image from the OpenShift Container Platform release image by running the following command:

```
$ CCO_IMAGE=$(oc adm release info --image-for='cloud-credential-operator'
$RELEASE_IMAGE -a ~/.pull-secret)
```



NOTE

Ensure that the architecture of the **\$RELEASE_IMAGE** matches the architecture of the environment in which you will use the **ccoctl** tool.

3. Extract the **ccoctl** binary from the CCO container image within the OpenShift Container Platform release image by running the following command:

```
$ oc image extract $CCO_IMAGE \
--file="/usr/bin/ccoctl.<rhel_version>" \
-a ~/.pull-secret
```

- 1 For **<rhel_version>**, specify the value that corresponds to the version of Red Hat Enterprise Linux (RHEL) that the host uses. If no value is specified, **ccoctl.rhel8** is used by default. The following values are valid:

- **rhel8**: Specify this value for hosts that use RHEL 8.
- **rhel9**: Specify this value for hosts that use RHEL 9.

4. Change the permissions to make **ccoctl** executable by running the following command:

```
$ chmod 775 ccoctl.<rhel_version>
```

Verification

- To verify that **ccoctl** is ready to use, display the help file by running the following command:

```
$ ccoctl --help
```

Output of `ccoctl --help`

OpenShift credentials provisioning tool

Usage:

```
ccoctl [command]
```

Available Commands:

```
aws      Manage credentials objects for AWS cloud
azure    Manage credentials objects for Azure
gcp      Manage credentials objects for Google cloud
help     Help about any command
ibmcloud Manage credentials objects for IBM Cloud
nutanix  Manage credentials objects for Nutanix
```

Flags:

```
-h, --help  help for ccoctl
```

Use "`ccoctl [command] --help`" for more information about a command.

7.5.10.2.2. Creating Azure resources with the Cloud Credential Operator utility

You can use the `ccoctl azure create-all` command to automate the creation of Azure resources.



NOTE

By default, `ccoctl` creates objects in the directory in which the commands are run. To create the objects in a different directory, use the `--output-dir` flag. This procedure uses `<path_to_ccoctl_output_dir>` to refer to this directory.

Prerequisites

You must have:

- Extracted and prepared the `ccoctl` binary.
- Access to your Microsoft Azure account by using the Azure CLI.

Procedure

1. Set a `$RELEASE_IMAGE` variable with the release image from your installation file by running the following command:


```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

2. Extract the list of `CredentialsRequest` objects from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
```

```
--included \1
--install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \2
--to=<path_to_directory_for_credentials_requests> \3
```

- 1 The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2 Specify the location of the **install-config.yaml** file.
- 3 Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.



NOTE

This command might take a few moments to run.


3. To enable the **ccoctl** utility to detect your Azure credentials automatically, log in to the Azure CLI by running the following command:

```
$ az login
```

4. Use the **ccoctl** tool to process all **CredentialsRequest** objects by running the following command:

```
$ ccoctl azure create-all \
--name=<azure_infra_name> \1
--output-dir=<ccoctl_output_dir> \2
--region=<azure_region> \3
--subscription-id=<azure_subscription_id> \4
--credentials-requests-dir=<path_to_credentials_requests_directory> \5
--dnszone-resource-group-name=<azure_dns_zone_resource_group_name> \6
--tenant-id=<azure_tenant_id> \7
```

- 1 Specify the user-defined name for all created Azure resources used for tracking.
- 2 Optional: Specify the directory in which you want the **ccoctl** utility to create objects. By default, the utility creates objects in the directory in which the commands are run.
- 3 Specify the Azure region in which cloud resources will be created.
- 4 Specify the Azure subscription ID to use.
- 5 Specify the directory containing the files for the component **CredentialsRequest** objects.
- 6 Specify the name of the resource group containing the cluster's base domain Azure DNS zone.
- 7 Specify the Azure tenant ID to use.



NOTE

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

To see additional optional parameters and explanations of how to use them, run the **azure create-all --help** command.

Verification

- To verify that the OpenShift Container Platform secrets are created, list the files in the **<path_to_ccoctl_output_dir>/manifests** directory:

```
$ ls <path_to_ccoctl_output_dir>/manifests
```

Example output

```
azure-ad-pod-identity-webhook-config.yaml
cluster-authentication-02-config.yaml
openshift-cloud-controller-manager-azure-cloud-credentials-credentials.yaml
openshift-cloud-network-config-controller-cloud-credentials-credentials.yaml
openshift-cluster-api-capz-manager-bootstrap-credentials-credentials.yaml
openshift-cluster-csi-drivers-azure-disk-credentials-credentials.yaml
openshift-cluster-csi-drivers-azure-file-credentials-credentials.yaml
openshift-image-registry-installer-cloud-credentials-credentials.yaml
openshift-ingress-operator-cloud-credentials-credentials.yaml
openshift-machine-api-azure-cloud-credentials-credentials.yaml
```

You can verify that the Microsoft Entra ID service accounts are created by querying Azure. For more information, refer to Azure documentation on listing Entra ID service accounts.

7.5.10.2.3. Incorporating the Cloud Credential Operator utility manifests

To implement short-term security credentials managed outside the cluster for individual components, you must move the manifest files that the Cloud Credential Operator utility (**ccoctl**) created to the correct directories for the installation program.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have configured the Cloud Credential Operator utility (**ccoctl**).
- You have created the cloud provider resources that are required for your cluster with the **ccoctl** utility.

Procedure

1. If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```

apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...

```

- If you used the **ccoctl** utility to create a new Azure resource group instead of using an existing resource group, modify the **resourceGroupName** parameter in the **install-config.yaml** as shown:

Sample configuration file snippet

```

apiVersion: v1
baseDomain: example.com
# ...
platform:
  azure:
    resourceGroupName: <azure_infra_name> 1
# ...

```

- This value must match the user-defined name for Azure resources that was specified with the **--name** argument of the **ccoctl azure create-all** command.

- If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

- Copy the manifests that the **ccoctl** utility generated to the **manifests** directory that the installation program created by running the following command:

```
$ cp /<path_to_ccoctl_output_dir>/manifests/* ./manifests/
```

- Copy the private key that the **ccoctl** utility generated in the **tls** directory to the installation directory by running the following command:

```
$ cp -a /<path_to_ccoctl_output_dir>/tls .
```

7.5.11. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.

- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have an Azure subscription ID and tenant ID.

Procedure

- Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

- 1 For **<installation_directory>**, specify the location of your customized **./install-config.yaml** file.
- 2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to **<installation_directory>/openshift_install.log**.



IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

7.5.12. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

Additional resources

- See [Accessing the web console](#) for more details about accessing and understanding the OpenShift Container Platform web console.

7.5.13. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service

7.5.14. Next steps

- [Customize your cluster](#).
- If necessary, you can [opt out of remote health reporting](#) .

7.6. INSTALLING A CLUSTER ON AZURE WITH NETWORK CUSTOMIZATIONS

In OpenShift Container Platform version 4.16, you can install a cluster with a customized network configuration on infrastructure that the installation program provisions on Microsoft Azure. By customizing your network configuration, your cluster can coexist with existing IP address allocations in your environment and integrate with existing MTU and VXLAN configurations.

You must set most of the network configuration parameters during installation, and you can modify only **kubeProxy** configuration parameters in a running cluster.

7.6.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [configured an Azure account](#) to host the cluster and determined the tested and validated region to deploy the cluster to.
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.
- If you use customer-managed encryption keys, you [prepared your Azure environment for encryption](#).

7.6.2. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.

- Access [Quay.io](https://quay.io) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.



IMPORTANT

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

7.6.3. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.



NOTE


You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.

**NOTE**

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.


2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.

**NOTE**

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```

**NOTE**

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

7.6.4. Obtaining the installation program

Before you install OpenShift Container Platform, download the installation file on the host you are using for installation.

Prerequisites

- You have a computer that runs Linux or macOS, with 500 MB of local disk space.

Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select your infrastructure provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.



IMPORTANT

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.



IMPORTANT

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.

4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#) . This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

7.6.5. Creating the installation configuration file

You can customize the OpenShift Container Platform cluster you install on Microsoft Azure.

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have an Azure subscription ID and tenant ID.
- If you are installing the cluster using a service principal, you have its application ID and password.
- If you are installing the cluster using a system-assigned managed identity, you have enabled it on the virtual machine that you will run the installation program from.
- If you are installing the cluster using a user-assigned managed identity, you have met these prerequisites:
 - You have its client ID.
 - You have assigned it to the virtual machine that you will run the installation program from.

Procedure

1. Optional: If you have run the installation program on this computer before, and want to use an alternative service principal or managed identity, go to the `~/.azure/` directory and delete the **osServicePrincipal.json** configuration file.
Deleting this file prevents the installation program from automatically reusing subscription and authentication values from a previous installation.

2. Create the **install-config.yaml** file.

- a. Change to the directory that contains the installation program and run the following command:

```

$ ./openshift-install create install-config --dir <installation_directory> 1

```

- 1** For **<installation_directory>**, specify the directory name to store the files that the installation program creates.

When specifying the directory:

- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
 - Use an empty directory. Some installation assets, such as bootstrap X.509 certificates, have short expiration intervals, therefore you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.
- b. At the prompts, provide the configuration details for your cloud:
 - i. Optional: Select an SSH key to use to access your cluster machines.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- ii. Select **azure** as the platform to target.
If the installation program cannot locate the **osServicePrincipal.json** configuration file from a previous installation, you are prompted for Azure subscription and authentication values.
- iii. Enter the following Azure parameter values for your subscription:
 - **azure subscription id** Enter the subscription ID to use for the cluster.
 - **azure tenant id** Enter the tenant ID.
- iv. Depending on the Azure identity you are using to deploy the cluster, do one of the following when prompted for the **azure service principal client id**
 - If you are using a service principal, enter its application ID.
 - If you are using a system-assigned managed identity, leave this value blank.
 - If you are using a user-assigned managed identity, specify its client ID.
- v. Depending on the Azure identity you are using to deploy the cluster, do one of the following when prompted for the **azure service principal client secret**
 - If you are using a service principal, enter its password.
 - If you are using a system-assigned managed identity, leave this value blank.
 - If you are using a user-assigned managed identity, leave this value blank.
- vi. Select the region to deploy the cluster to.
- vii. Select the base domain to deploy the cluster to. The base domain corresponds to the Azure DNS Zone that you created for your cluster.
- viii. Enter a descriptive name for your cluster.



IMPORTANT

All Azure resources that are available through public endpoints are subject to resource name restrictions, and you cannot create resources that use certain terms. For a list of terms that Azure restricts, see [Resolve reserved resource name errors](#) in the Azure documentation.

3. Modify the **install-config.yaml** file. You can find more information about the available parameters in the "Installation configuration parameters" section.
4. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

If previously not detected, the installation program creates an **osServicePrincipal.json** configuration file and stores this file in the `~/.azure/` directory on your computer. This ensures that the installation program can load the profile when it is creating an OpenShift Container Platform cluster on the target

platform.

Additional resources

- [Installation configuration parameters for Azure](#)

7.6.5.1. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 7.2. Minimum resource requirements

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Bootstrap	RHCOS	4	16 GB	100 GB	300
Control plane	RHCOS	4	16 GB	100 GB	300
Compute	RHCOS, RHEL 8.6 and later [3]	2	8 GB	100 GB	300

1. One vCPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio: $(\text{threads per core} \times \text{cores}) \times \text{sockets} = \text{vCPUs}$.
2. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes which require a 10 ms p99 fsync duration. Note that on many cloud platforms, storage size and IOPS scale together, so you might need to over-allocate storage volume to obtain sufficient performance.
3. As with all user-provisioned installations, if you choose to use RHEL compute machines in your cluster, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Use of RHEL 7 compute machines is deprecated and has been removed in OpenShift Container Platform 4.10 and later.




NOTE

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).



IMPORTANT

You are required to use Azure virtual machines that have the **premiumIO** parameter set to **true**.

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

7.6.5.2. Tested instance types for Azure

The following Microsoft Azure instance types have been tested with OpenShift Container Platform.

Example 7.27. Machine types based on 64-bit x86 architecture

- **c4.***
- **c5.***
- **c5a.***
- **i3.***
- **m4.***
- **m5.***
- **m5a.***
- **m6a.***
- **m6i.***
- **r4.***
- **r5.***

- **r5a.***
- **r6i.***
- **t3.***
- **t3a.***

7.6.5.3. Tested instance types for Azure on 64-bit ARM infrastructures

The following Microsoft Azure ARM64 instance types have been tested with OpenShift Container Platform.


Example 7.28. Machine types based on 64-bit ARM architecture

- **c6g.***
- **m6g.***

7.6.5.4. Enabling trusted launch for Azure VMs

You can enable two trusted launch features when installing your cluster on Azure: [secure boot](#) and [virtualized Trusted Platform Modules](#).

See the Azure documentation about [virtual machine sizes](#) to learn what sizes of virtual machines support these features.



IMPORTANT

Trusted launch is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

Prerequisites

- You have created an **install-config.yaml** file.

Procedure

- Use a text editor to edit the **install-config.yaml** file prior to deploying your cluster and add the following stanza:

```
controlPlane: ❶
platform:
  azure:
    settings:
      securityType: TrustedLaunch ❷
```

```

trustedLaunch:
  uefiSettings:
    secureBoot: Enabled 3
    virtualizedTrustedPlatformModule: Enabled 4

```

- 1** Specify **controlPlane.platform.azure** or **compute.platform.azure** to enable trusted launch on only control plane or compute nodes respectively. Specify **platform.azure.defaultMachinePlatform** to enable trusted launch on all nodes.
- 2** Enable trusted launch features.
- 3** Enable secure boot. For more information, see the Azure documentation about [secure boot](#).
- 4** Enable the virtualized Trusted Platform Module. For more information, see the Azure documentation about [virtualized Trusted Platform Modules](#).

7.6.5.5. Enabling confidential VMs

You can enable confidential VMs when installing your cluster. You can enable confidential VMs for compute nodes, control plane nodes, or all nodes.



IMPORTANT

Using confidential VMs is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

You can use confidential VMs with the following VM sizes:

- DCasv5-series
- DCadsv5-series
- ECasv5-series
- ECadsv5-series



IMPORTANT

Confidential VMs are currently not supported on 64-bit ARM architectures.

Prerequisites

- You have created an **install-config.yaml** file.

Procedure

- Use a text editor to edit the **install-config.yaml** file prior to deploying your cluster and add the following stanza:

```
controlPlane: ❶
  platform:
    azure:
      settings:
        securityType: ConfidentialVM ❷
        confidentialVM:
          uefiSettings:
            secureBoot: Enabled ❸
            virtualizedTrustedPlatformModule: Enabled ❹
      osDisk:
        securityProfile:
          securityEncryptionType: VMGuestStateOnly ❺
```

- ❶ Specify **controlPlane.platform.azure** or **compute.platform.azure** to deploy confidential VMs on only control plane or compute nodes respectively. Specify **platform.azure.defaultMachinePlatform** to deploy confidential VMs on all nodes.
- ❷ Enable confidential VMs.
- ❸ Enable secure boot. For more information, see the Azure documentation about [secure boot](#).
- ❹ Enable the virtualized Trusted Platform Module. For more information, see the Azure documentation about [virtualized Trusted Platform Modules](#).
- ❺ Specify **VMGuestStateOnly** to encrypt the VM guest state.

7.6.5.6. Sample customized install-config.yaml file for Azure

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.



IMPORTANT

This sample YAML file is provided for reference only. You must obtain your **install-config.yaml** file by using the installation program and modify it.

```
apiVersion: v1
baseDomain: example.com ❶
controlPlane: ❷
  hyperthreading: Enabled ❸ ❹
  name: master
  platform:
    azure:
      encryptionAtHost: true
      ultraSSDCapability: Enabled
      osDisk:
        diskSizeGB: 1024 ❺
        diskType: Premium_LRS
        diskEncryptionSet:
```

```
    resourceGroup: disk_encryption_set_resource_group
    name: disk_encryption_set_name
    subscriptionId: secondary_subscription_id
  osImage:
    publisher: example_publisher_name
    offer: example_image_offer
    sku: example_offer_sku
    version: example_image_version
    type: Standard_D8s_v3
  replicas: 3
compute: 6
- hyperthreading: Enabled 7
  name: worker
  platform:
    azure:
      ultraSSDCapability: Enabled
      type: Standard_D2s_v3
      encryptionAtHost: true
    osDisk:
      diskSizeGB: 512 8
      diskType: Standard_LRS
      diskEncryptionSet:
        resourceGroup: disk_encryption_set_resource_group
        name: disk_encryption_set_name
        subscriptionId: secondary_subscription_id
    osImage:
      publisher: example_publisher_name
      offer: example_image_offer
      sku: example_offer_sku
      version: example_image_version
  zones: 9
  - "1"
  - "2"
  - "3"
  replicas: 5
metadata:
  name: test-cluster 10
networking: 11
  clusterNetwork:
    - cidr: 10.128.0.0/14
      hostPrefix: 23
  machineNetwork:
    - cidr: 10.0.0.0/16
  networkType: OVNKubernetes 12
  serviceNetwork:
    - 172.30.0.0/16
platform:
  azure:
    defaultMachinePlatform:
      osImage: 13
      publisher: example_publisher_name
      offer: example_image_offer
      sku: example_offer_sku
      version: example_image_version
      ultraSSDCapability: Enabled
```

```

baseDomainResourceGroupName: resource_group 14
region: centralus 15
resourceGroupName: existing_resource_group 16
outboundType: Loadbalancer
cloudName: AzurePublicCloud
pullSecret: '{"auths": ...}' 17
fips: false 18
sshKey: ssh-ed25519 AAAA... 19

```

- 1 10 15 17 Required. The installation program prompts you for this value.
- 2 6 11 If you do not provide these parameters and values, the installation program provides the default value.
- 3 7 The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.
- 4 Whether to enable or disable simultaneous multithreading, or **hypertreading**. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores. You can disable it by setting the parameter value to **Disabled**. If you disable simultaneous multithreading in some cluster machines, you must disable it in all cluster machines.



IMPORTANT

If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance. Use larger virtual machine types, such as **Standard_D8s_v3**, for your machines if you disable simultaneous multithreading.

- 5 8 You can specify the size of the disk to use in GB. Minimum recommendation for control plane nodes is 1024 GB.
- 9 Specify a list of zones to deploy your machines to. For high availability, specify at least two zones.
- 12 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 13 Optional: A custom Red Hat Enterprise Linux CoreOS (RHCOS) image that should be used to boot control plane and compute machines. The **publisher**, **offer**, **sku**, and **version** parameters under **platform.azure.defaultMachinePlatform.osImage** apply to both control plane and compute machines. If the parameters under **controlPlane.platform.azure.osImage** or **compute.platform.azure.osImage** are set, they override the **platform.azure.defaultMachinePlatform.osImage** parameters.
- 14 Specify the name of the resource group that contains the DNS zone for your base domain.
- 16 Specify the name of an already existing resource group to install your cluster to. If undefined, a new resource group is created for the cluster.
- 18 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.



IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 19** You can optionally provide the **sshKey** value that you use to access the machines in your cluster.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

7.6.5.7. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
```



```

proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5


```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.



NOTE

The installation program does not support the proxy **readinessEndpoints** field.



NOTE

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.

**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

7.6.6. Network configuration phases


There are two phases prior to OpenShift Container Platform installation where you can customize the network configuration.

Phase 1

You can customize the following network-related fields in the **install-config.yaml** file before you create the manifest files:

- **networking.networkType**
- **networking.clusterNetwork**
- **networking.serviceNetwork**
- **networking.machineNetwork**

For more information on these fields, refer to *Installation configuration parameters*.

**NOTE**

Set the **networking.machineNetwork** to match the CIDR that the preferred NIC resides in.

**IMPORTANT**

The CIDR range **172.17.0.0/16** is reserved by libVirt. You cannot use this range or any range that overlaps with this range for any networks in your cluster.

Phase 2

After creating the manifest files by running **openshift-install create manifests**, you can define a customized Cluster Network Operator manifest with only the fields you want to modify. You can use the manifest to specify advanced network configuration.

You cannot override the values specified in phase 1 in the **install-config.yaml** file during phase 2. However, you can further customize the network plugin during phase 2.

7.6.7. Specifying advanced network configuration

You can use advanced network configuration for your network plugin to integrate your cluster into your existing network environment. You can specify advanced network configuration only before you install the cluster.

**IMPORTANT**

Customizing your network configuration by modifying the OpenShift Container Platform manifest files created by the installation program is not supported. Applying a manifest file that you create, as in the following procedure, is supported.

Prerequisites

- You have created the **install-config.yaml** file and completed any modifications to it.

Procedure

- Change to the directory that contains the installation program and create the manifests:

```
$ ./openshift-install create manifests --dir <installation_directory> 1
```

- 1** **<installation_directory>** specifies the name of the directory that contains the **install-config.yaml** file for your cluster.

- Create a stub manifest file for the advanced network configuration that is named **cluster-network-03-config.yml** in the **<installation_directory>/manifests/** directory:

```
apiVersion: operator.openshift.io/v1
kind: Network
metadata:
  name: cluster
spec:
```

- Specify the advanced network configuration for your cluster in the **cluster-network-03-config.yml** file, such as in the following example:

Enable IPsec for the OVN-Kubernetes network provider

```
apiVersion: operator.openshift.io/v1
kind: Network
metadata:
  name: cluster
spec:
  defaultNetwork:
    ovnKubernetesConfig:
      ipsecConfig:
        mode: Full
```

- Optional: Back up the **manifests/cluster-network-03-config.yml** file. The installation program consumes the **manifests/** directory when you create the Ignition config files.

7.6.8. Cluster Network Operator configuration

The configuration for the cluster network is specified as part of the Cluster Network Operator (CNO) configuration and stored in a custom resource (CR) object that is named **cluster**. The CR specifies the fields for the **Network** API in the **operator.openshift.io** API group.

The CNO configuration inherits the following fields during cluster installation from the **Network** API in the **Network.config.openshift.io** API group:

clusterNetwork

IP address pools from which pod IP addresses are allocated.

serviceNetwork

IP address pool for services.

defaultNetwork.type

Cluster network plugin. **OVNKubernetes** is the only supported plugin during installation.

You can specify the cluster network plugin configuration for your cluster by setting the fields for the **defaultNetwork** object in the CNO object named **cluster**.

7.6.8.1. Cluster Network Operator configuration object

The fields for the Cluster Network Operator (CNO) are described in the following table:


Table 7.3. Cluster Network Operator configuration object

Field	Type	Description
metadata.name	string	The name of the CNO object. This name is always cluster .
spec.clusterNetwork	array	<p>A list specifying the blocks of IP addresses from which pod IP addresses are allocated and the subnet prefix length assigned to each individual node in the cluster. For example:</p> <pre>spec: clusterNetwork: - cidr: 10.128.0.0/19 hostPrefix: 23 - cidr: 10.128.32.0/19 hostPrefix: 23</pre>
spec.serviceNetwork	array	<p>A block of IP addresses for services. The OpenShift SDN and OVN-Kubernetes network plugins support only a single IP address block for the service network. For example:</p> <pre>spec: serviceNetwork: - 172.30.0.0/14</pre> <p>You can customize this field only in the install-config.yaml file before you create the manifests. The value is read-only in the manifest file.</p>
spec.defaultNetwork	object	Configures the network plugin for the cluster network.
spec.kubeProxyConfig	object	The fields for this object specify the kube-proxy configuration. If you are using the OVN-Kubernetes cluster network plugin, the kube-proxy configuration has no effect.

defaultNetwork object configuration

The values for the **defaultNetwork** object are defined in the following table:

Table 7.4. defaultNetwork object

Field	Type	Description
type	string	<p>OVNKubernetes. The Red Hat OpenShift Networking network plugin is selected during installation. This value cannot be changed after cluster installation.</p> <div style="display: flex; align-items: flex-start;">  <div> <p>NOTE</p> <p>OpenShift Container Platform uses the OVN-Kubernetes network plugin by default. OpenShift SDN is no longer available as an installation choice for new clusters.</p> </div> </div>
ovnKubernetesConfig	object	This object is only valid for the OVN-Kubernetes network plugin.

Configuration for the OVN-Kubernetes network plugin

The following table describes the configuration fields for the OVN-Kubernetes network plugin:

Table 7.5. ovnKubernetesConfig object

Field	Type	Description
mtu	integer	<p>The maximum transmission unit (MTU) for the Geneve (Generic Network Virtualization Encapsulation) overlay network. This is detected automatically based on the MTU of the primary network interface. You do not normally need to override the detected MTU.</p> <p>If the auto-detected value is not what you expect it to be, confirm that the MTU on the primary network interface on your nodes is correct. You cannot use this option to change the MTU value of the primary network interface on the nodes.</p> <p>If your cluster requires different MTU values for different nodes, you must set this value to 100 less than the lowest MTU value in your cluster. For example, if some nodes in your cluster have an MTU of 9001, and some have an MTU of 1500, you must set this value to 1400.</p>
genevePort	integer	The port to use for all Geneve packets. The default value is 6081 . This value cannot be changed after cluster installation.
ipsecConfig	object	Specify a configuration object for customizing the IPsec configuration.
ipv4	object	Specifies a configuration object for IPv4 settings.
ipv6	object	Specifies a configuration object for IPv6 settings.


Field	Type	Description
policyAuditConfig	object	Specify a configuration object for customizing network policy audit logging. If unset, the defaults audit log settings are used.
gatewayConfig	object	Optional: Specify a configuration object for customizing how egress traffic is sent to the node gateway.  <p>NOTE</p> <p>While migrating egress traffic, you can expect some disruption to workloads and service traffic until the Cluster Network Operator (CNO) successfully rolls out the changes.</p>

Table 7.6. `ovnKubernetesConfig.ipv4` object

Field	Type	Description
internalTransitSwitchSubnet	string	If your existing network infrastructure overlaps with the 100.88.0.0/16 IPv4 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. The subnet for the distributed transit switch that enables east-west traffic. This subnet cannot overlap with any other subnets used by OVN-Kubernetes or on the host itself. It must be large enough to accommodate one IP address per node in your cluster. The default value is 100.88.0.0/16 .
internalJoinSubnet	string	If your existing network infrastructure overlaps with the 100.64.0.0/16 IPv4 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. You must ensure that the IP address range does not overlap with any other subnet used by your OpenShift Container Platform installation. The IP address range must be larger than the maximum number of nodes that can be added to the cluster. For example, if the clusterNetwork.cidr value is 10.128.0.0/14 and the clusterNetwork.hostPrefix value is /23 , then the maximum number of nodes is 2⁽²³⁻¹⁴⁾=512 . The default value is 100.64.0.0/16 .

Table 7.7. `ovnKubernetesConfig.ipv6` object

Field	Type	Description
-------	------	-------------

Field	Type	Description
internalTransitSwitchSubnet	string	<p>If your existing network infrastructure overlaps with the fd97::/64 IPv6 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. The subnet for the distributed transit switch that enables east-west traffic. This subnet cannot overlap with any other subnets used by OVN-Kubernetes or on the host itself. It must be large enough to accommodate one IP address per node in your cluster.</p> <p>The default value is fd97::/64.</p>
internalJoinSubnet	string	<p>If your existing network infrastructure overlaps with the fd98::/64 IPv6 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. You must ensure that the IP address range does not overlap with any other subnet used by your OpenShift Container Platform installation. The IP address range must be larger than the maximum number of nodes that can be added to the cluster.</p> <p>The default value is fd98::/64.</p>

Table 7.8. `policyAuditConfig` object

Field	Type	Description
rateLimit	integer	The maximum number of messages to generate every second per node. The default value is 20 messages per second.
maxFileSize	integer	The maximum size for the audit log in bytes. The default value is 50000000 or 50 MB.
maxLogFiles	integer	The maximum number of log files that are retained.
destination	string	<p>One of the following additional audit log targets:</p> <p>libc The libc syslog() function of the journald process on the host.</p> <p>udp:<host>:<port> A syslog server. Replace <host>:<port> with the host and port of the syslog server.</p> <p>unix:<file> A Unix Domain Socket file specified by <file>.</p> <p>null Do not send the audit logs to any additional target.</p>
syslogFacility	string	The syslog facility, such as kern , as defined by RFC5424. The default value is local0 .

Table 7.9. gatewayConfig object

Field	Type	Description
routingViaHost	boolean	<p>Set this field to true to send egress traffic from pods to the host networking stack. For highly-specialized installations and applications that rely on manually configured routes in the kernel routing table, you might want to route egress traffic to the host networking stack. By default, egress traffic is processed in OVN to exit the cluster and is not affected by specialized routes in the kernel routing table. The default value is false.</p> <p>This field has an interaction with the Open vSwitch hardware offloading feature. If you set this field to true, you do not receive the performance benefits of the offloading because egress traffic is processed by the host networking stack.</p>
ipForwarding	object	<p>You can control IP forwarding for all traffic on OVN-Kubernetes managed interfaces by using the ipForwarding specification in the Network resource. Specify Restricted to only allow IP forwarding for Kubernetes related traffic. Specify Global to allow forwarding of all IP traffic. For new installations, the default is Restricted. For updates to OpenShift Container Platform 4.14 or later, the default is Global.</p>
ipv4	object	<p>Optional: Specify an object to configure the internal OVN-Kubernetes masquerade address for host to service traffic for IPv4 addresses.</p>
ipv6	object	<p>Optional: Specify an object to configure the internal OVN-Kubernetes masquerade address for host to service traffic for IPv6 addresses.</p>

Table 7.10. gatewayConfig.ipv4 object

Field	Type	Description
internalMasqueradeSubnet	string	<p>The masquerade IPv4 addresses that are used internally to enable host to service traffic. The host is configured with these IP addresses as well as the shared gateway bridge interface. The default value is 169.254.169.0/29.</p>

Table 7.11. gatewayConfig.ipv6 object


Field	Type	Description
internalMasqueradeSubnet	string	<p>The masquerade IPv6 addresses that are used internally to enable host to service traffic. The host is configured with these IP addresses as well as the shared gateway bridge interface. The default value is fd69::/125.</p>

Table 7.12. `ipsecConfig` object

Field	Type	Description
<code>mode</code>	<code>string</code>	<p>Specifies the behavior of the IPsec implementation. Must be one of the following values:</p> <ul style="list-style-type: none"> ● Disabled: IPsec is not enabled on cluster nodes. ● External: IPsec is enabled for network traffic with external hosts. ● Full: IPsec is enabled for pod traffic and network traffic with external hosts.

Example OVN-Kubernetes configuration with IPsec enabled

```
defaultNetwork:
  type: OVNKubernetes
  ovnKubernetesConfig:
    mtu: 1400
    genevePort: 6081
    ipsecConfig:
      mode: Full
```




IMPORTANT

Using OVNKubernetes can lead to a stack exhaustion problem on IBM Power®.

kubeProxyConfig object configuration (OpenShiftSDN container network interface only)

The values for the `kubeProxyConfig` object are defined in the following table:

Table 7.13. `kubeProxyConfig` object

Field	Type	Description
<code>iptablesSyncPeriod</code>	<code>string</code>	<p>The refresh period for <code>iptables</code> rules. The default value is 30s. Valid suffixes include s, m, and h and are described in the Go time package documentation.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p>NOTE</p> <p>Because of performance improvements introduced in OpenShift Container Platform 4.3 and greater, adjusting the <code>iptablesSyncPeriod</code> parameter is no longer necessary.</p> </div> </div>

Field	Type	Description
proxyArguments.iptables-min-sync-period	array	<p>The minimum duration before refreshing iptables rules. This field ensures that the refresh does not happen too frequently. Valid suffixes include s, m, and h and are described in the Go time package. The default value is:</p> <pre>kubeProxyConfig: proxyArguments: iptables-min-sync-period: - 0s</pre>

7.6.9. Configuring hybrid networking with OVN-Kubernetes

You can configure your cluster to use hybrid networking with the OVN-Kubernetes network plugin. This allows a hybrid cluster that supports different node networking configurations.

Prerequisites

- You defined **OVNKubernetes** for the **networking.networkType** parameter in the **install-config.yaml** file. See the installation documentation for configuring OpenShift Container Platform network customizations on your chosen cloud provider for more information.

Procedure

- Change to the directory that contains the installation program and create the manifests:

```
$ ./openshift-install create manifests --dir <installation_directory>
```

where:

<installation_directory>

Specifies the name of the directory that contains the **install-config.yaml** file for your cluster.

- Create a stub manifest file for the advanced network configuration that is named **cluster-network-03-config.yml** in the **<installation_directory>/manifests/** directory:

```
$ cat <<EOF > <installation_directory>/manifests/cluster-network-03-config.yml
apiVersion: operator.openshift.io/v1
kind: Network
metadata:
  name: cluster
spec:
EOF
```

where:

<installation_directory>

Specifies the directory name that contains the **manifests/** directory for your cluster.

3. Open the **cluster-network-03-config.yml** file in an editor and configure OVN-Kubernetes with hybrid networking, such as in the following example:

Specify a hybrid networking configuration

```
apiVersion: operator.openshift.io/v1
kind: Network
metadata:
  name: cluster
spec:
  defaultNetwork:
    ovnKubernetesConfig:
      hybridOverlayConfig:
        hybridClusterNetwork: 1
        - cidr: 10.132.0.0/14
        hostPrefix: 23
```

- 1** Specify the CIDR configuration used for nodes on the additional overlay network. The **hybridClusterNetwork** CIDR cannot overlap with the **clusterNetwork** CIDR.

4. Save the **cluster-network-03-config.yml** file and quit the text editor.
5. Optional: Back up the **manifests/cluster-network-03-config.yml** file. The installation program deletes the **manifests/** directory when creating the cluster.

Additional resources

- For more details about Accelerated Networking, see [Accelerated Networking for Microsoft Azure VMs](#).

7.6.10. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.



IMPORTANT

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.

4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**.
To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:


```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.

**NOTE**

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your PATH.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

7.6.11. Alternatives to storing administrator-level secrets in the kube-system project

By default, administrator secrets are stored in the **kube-system** project. If you configured the **credentialsMode** parameter in the **install-config.yaml** file to **Manual**, you must use one of the following alternatives:

- To manage long-term cloud credentials manually, follow the procedure in [Manually creating long-term credentials](#).
- To implement short-term credentials that are managed outside the cluster for individual components, follow the procedures in [Configuring an Azure cluster to use short-term credentials](#).

7.6.11.1. Manually creating long-term credentials

The Cloud Credential Operator (CCO) can be put into manual mode prior to installation in environments where the cloud identity and access management (IAM) APIs are not reachable, or the administrator prefers not to store an administrator-level credential secret in the cluster **kube-system** namespace.

Procedure

1. If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```
apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...
```

2. If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

- Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

- Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included 1 \
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml 2 \
  --to=<path_to_directory_for_credentials_requests> 3
```

- The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- Specify the location of the **install-config.yaml** file.
- Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

This command creates a YAML file for each **CredentialsRequest** object.

Sample CredentialsRequest object

```
apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: AzureProviderSpec
    roleBindings:
      - role: Contributor
    ...
```

- Create YAML files for secrets in the **openshift-install** manifests directory that you generated previously. The secrets must be stored using the namespace and secret name defined in the **spec.secretRef** for each **CredentialsRequest** object.

Sample CredentialsRequest object with secrets

```
apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
```

```

namespace: openshift-cloud-credential-operator
...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: AzureProviderSpec
    roleBindings:
      - role: Contributor
    ...
  secretRef:
    name: <component_secret>
    namespace: <component_namespace>
  ...


```

Sample Secret object

```

apiVersion: v1
kind: Secret
metadata:
  name: <component_secret>
  namespace: <component_namespace>
data:
  azure_subscription_id: <base64_encoded_azure_subscription_id>
  azure_client_id: <base64_encoded_azure_client_id>
  azure_client_secret: <base64_encoded_azure_client_secret>
  azure_tenant_id: <base64_encoded_azure_tenant_id>
  azure_resource_prefix: <base64_encoded_azure_resource_prefix>
  azure_resourcegroup: <base64_encoded_azure_resourcegroup>
  azure_region: <base64_encoded_azure_region>

```



IMPORTANT


Before upgrading a cluster that uses manually maintained credentials, you must ensure that the CCO is in an upgradeable state.

7.6.11.2. Configuring an Azure cluster to use short-term credentials

To install a cluster that uses Microsoft Entra Workload ID, you must configure the Cloud Credential Operator utility and create the required Azure resources for your cluster.

7.6.11.2.1. Configuring the Cloud Credential Operator utility

To create and manage cloud credentials from outside of the cluster when the Cloud Credential Operator (CCO) is operating in manual mode, extract and prepare the CCO utility (**ccoctl**) binary.



NOTE

The **ccoctl** utility is a Linux binary that must run in a Linux environment.

Prerequisites

- You have access to an OpenShift Container Platform account with cluster administrator access.
- You have installed the OpenShift CLI (**oc**).

- You have created a global Microsoft Azure account for the **ccoctl** utility to use with the following permissions:

Example 7.29. Required Azure permissions

- Microsoft.Resources/subscriptions/resourceGroups/read
- Microsoft.Resources/subscriptions/resourceGroups/write
- Microsoft.Resources/subscriptions/resourceGroups/delete
- Microsoft.Authorization/roleAssignments/read
- Microsoft.Authorization/roleAssignments/delete
- Microsoft.Authorization/roleAssignments/write
- Microsoft.Authorization/roleDefinitions/read
- Microsoft.Authorization/roleDefinitions/write
- Microsoft.Authorization/roleDefinitions/delete
- Microsoft.Storage/storageAccounts/listkeys/action
- Microsoft.Storage/storageAccounts/delete
- Microsoft.Storage/storageAccounts/read
- Microsoft.Storage/storageAccounts/write
- Microsoft.Storage/storageAccounts/blobServices/containers/write
- Microsoft.Storage/storageAccounts/blobServices/containers/delete
- Microsoft.Storage/storageAccounts/blobServices/containers/read
- Microsoft.ManagedIdentity/userAssignedIdentities/delete
- Microsoft.ManagedIdentity/userAssignedIdentities/read
- Microsoft.ManagedIdentity/userAssignedIdentities/write
- Microsoft.ManagedIdentity/userAssignedIdentities/federatedIdentityCredentials/read
- Microsoft.ManagedIdentity/userAssignedIdentities/federatedIdentityCredentials/write
- Microsoft.ManagedIdentity/userAssignedIdentities/federatedIdentityCredentials/delete
- Microsoft.Storage/register/action
- Microsoft.ManagedIdentity/register/action


Procedure

1. Set a variable for the OpenShift Container Platform release image by running the following command:


```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

- Obtain the CCO container image from the OpenShift Container Platform release image by running the following command:

```
$ CCO_IMAGE=$(oc adm release info --image-for='cloud-credential-operator'
$RELEASE_IMAGE -a ~/.pull-secret)
```



NOTE

Ensure that the architecture of the **\$RELEASE_IMAGE** matches the architecture of the environment in which you will use the **ccoctl** tool.

- Extract the **ccoctl** binary from the CCO container image within the OpenShift Container Platform release image by running the following command:

```
$ oc image extract $CCO_IMAGE \
--file="/usr/bin/ccoctl.<rhel_version>" \
-a ~/.pull-secret
```

- For **<rhel_version>**, specify the value that corresponds to the version of Red Hat Enterprise Linux (RHEL) that the host uses. If no value is specified, **ccoctl.rhel8** is used by default. The following values are valid:

- rhel8**: Specify this value for hosts that use RHEL 8.
- rhel9**: Specify this value for hosts that use RHEL 9.

- Change the permissions to make **ccoctl** executable by running the following command:

```
$ chmod 775 ccoctl.<rhel_version>
```

Verification

- To verify that **ccoctl** is ready to use, display the help file by running the following command:

```
$ ccoctl --help
```

Output of **ccoctl --help**

```
OpenShift credentials provisioning tool
```

```
Usage:
ccoctl [command]
```

```
Available Commands:
```

```
aws      Manage credentials objects for AWS cloud
azure    Manage credentials objects for Azure
gcp      Manage credentials objects for Google cloud
help     Help about any command
ibmcloud Manage credentials objects for IBM Cloud
nutanix  Manage credentials objects for Nutanix
```

Flags:

-h, --help help for ccoctl

Use "ccoctl [command] --help" for more information about a command.

7.6.11.2.2. Creating Azure resources with the Cloud Credential Operator utility

You can use the **ccoctl azure create-all** command to automate the creation of Azure resources.

**NOTE**

By default, **ccoctl** creates objects in the directory in which the commands are run. To create the objects in a different directory, use the **--output-dir** flag. This procedure uses **<path_to_ccoctl_output_dir>** to refer to this directory.

Prerequisites

You must have:

- Extracted and prepared the **ccoctl** binary.
- Access to your Microsoft Azure account by using the Azure CLI.

Procedure

1. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

2. Extract the list of **CredentialsRequest** objects from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included 1 \
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml 2 \
  --to=<path_to_directory_for_credentials_requests> 3
```

- 1** The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2** Specify the location of the **install-config.yaml** file.
- 3** Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

**NOTE**

This command might take a few moments to run.

- To enable the **ccoctl** utility to detect your Azure credentials automatically, log in to the Azure CLI by running the following command:

```
$ az login
```

- Use the **ccoctl** tool to process all **CredentialsRequest** objects by running the following command:

```
$ ccoctl azure create-all \
  --name=<azure_infra_name> \ 1
  --output-dir=<ccoctl_output_dir> \ 2
  --region=<azure_region> \ 3
  --subscription-id=<azure_subscription_id> \ 4
  --credentials-requests-dir=<path_to_credentials_requests_directory> \ 5
  --dnszone-resource-group-name=<azure_dns_zone_resource_group_name> \ 6
  --tenant-id=<azure_tenant_id> \ 7
```

- Specify the user-defined name for all created Azure resources used for tracking.
- Optional: Specify the directory in which you want the **ccoctl** utility to create objects. By default, the utility creates objects in the directory in which the commands are run.
- Specify the Azure region in which cloud resources will be created.
- Specify the Azure subscription ID to use.
- Specify the directory containing the files for the component **CredentialsRequest** objects.
- Specify the name of the resource group containing the cluster's base domain Azure DNS zone.
- Specify the Azure tenant ID to use.



NOTE

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

To see additional optional parameters and explanations of how to use them, run the **azure create-all --help** command.

Verification

- To verify that the OpenShift Container Platform secrets are created, list the files in the **<path_to_ccoctl_output_dir>/manifests** directory:

```
$ ls <path_to_ccoctl_output_dir>/manifests
```

Example output

```
azure-ad-pod-identity-webhook-config.yaml
cluster-authentication-02-config.yaml
```

```

openshift-cloud-controller-manager-azure-cloud-credentials-credentials.yaml
openshift-cloud-network-config-controller-cloud-credentials-credentials.yaml
openshift-cluster-api-capz-manager-bootstrap-credentials-credentials.yaml
openshift-cluster-csi-drivers-azure-disk-credentials-credentials.yaml
openshift-cluster-csi-drivers-azure-file-credentials-credentials.yaml
openshift-image-registry-installer-cloud-credentials-credentials.yaml
openshift-ingress-operator-cloud-credentials-credentials.yaml
openshift-machine-api-azure-cloud-credentials-credentials.yaml

```

You can verify that the Microsoft Entra ID service accounts are created by querying Azure. For more information, refer to Azure documentation on listing Entra ID service accounts.

7.6.11.2.3. Incorporating the Cloud Credential Operator utility manifests

To implement short-term security credentials managed outside the cluster for individual components, you must move the manifest files that the Cloud Credential Operator utility (**ccoctl**) created to the correct directories for the installation program.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have configured the Cloud Credential Operator utility (**ccoctl**).
- You have created the cloud provider resources that are required for your cluster with the **ccoctl** utility.

Procedure

1. If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```

apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...

```

2. If you used the **ccoctl** utility to create a new Azure resource group instead of using an existing resource group, modify the **resourceGroupName** parameter in the **install-config.yaml** as shown:

Sample configuration file snippet

```

apiVersion: v1
baseDomain: example.com
# ...
platform:
  azure:
    resourceGroupName: <azure_infra_name> 1
# ...

```

- 1 This value must match the user-defined name for Azure resources that was specified with the **--name** argument of the **ccoctl azure create-all** command.
3. If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

4. Copy the manifests that the **ccoctl** utility generated to the **manifests** directory that the installation program created by running the following command:

```
$ cp /<path_to_ccoctl_output_dir>/manifests/* ./manifests/
```

5. Copy the private key that the **ccoctl** utility generated in the **tls** directory to the installation directory by running the following command:

```
$ cp -a /<path_to_ccoctl_output_dir>/tls .
```

7.6.12. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have an Azure subscription ID and tenant ID.

Procedure

- Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

1 For **<installation_directory>**, specify the location of your customized **./install-config.yaml** file.

2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to `<installation_directory>/openshift_install.log`.



IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

7.6.13. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

Additional resources

- See [Accessing the web console](#) for more details about accessing and understanding the OpenShift Container Platform web console.

7.6.14. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service

7.6.15. Next steps

- [Customize your cluster](#).
- If necessary, you can [opt out of remote health reporting](#) .

7.7. INSTALLING A CLUSTER ON AZURE INTO AN EXISTING VNET

In OpenShift Container Platform version 4.16, you can install a cluster into an existing Azure Virtual Network (VNet) on Microsoft Azure. The installation program provisions the rest of the required infrastructure, which you can further customize. To customize the installation, you modify parameters in the **install-config.yaml** file before you install the cluster.

7.7.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.

- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [configured an Azure account](#) to host the cluster and determined the tested and validated region to deploy the cluster to.
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.
- If you use customer-managed encryption keys, you [prepared your Azure environment for encryption](#).

7.7.2. About reusing a VNet for your OpenShift Container Platform cluster

In OpenShift Container Platform 4.16, you can deploy a cluster into an existing Azure Virtual Network (VNet) in Microsoft Azure. If you do, you must also use existing subnets within the VNet and routing rules.

By deploying OpenShift Container Platform into an existing Azure VNet, you might be able to avoid service limit constraints in new accounts or more easily abide by the operational constraints that your company's guidelines set. This is a good option to use if you cannot obtain the infrastructure creation permissions that are required to create the VNet.

7.7.2.1. Requirements for using your VNet

When you deploy a cluster by using an existing VNet, you must perform additional network configuration before you install the cluster. In installer-provisioned infrastructure clusters, the installer usually creates the following components, but it does not create them when you install into an existing VNet:

- Subnets
- Route tables
- VNets
- Network Security Groups



NOTE

The installation program requires that you use the cloud-provided DNS server. Using a custom DNS server is not supported and causes the installation to fail.

If you use a custom VNet, you must correctly configure it and its subnets for the installation program and the cluster to use. The installation program cannot subdivide network ranges for the cluster to use, set route tables for the subnets, or set VNet options like DHCP, so you must do so before you install the cluster.

The cluster must be able to access the resource group that contains the existing VNet and subnets. While all of the resources that the cluster creates are placed in a separate resource group that it creates, some network resources are used from a separate group. Some cluster Operators must be able to access resources in both resource groups. For example, the Machine API controller attaches NICs for the virtual machines that it creates to subnets from the networking resource group.

Your VNet must meet the following characteristics:

- The VNet's CIDR block must contain the **Networking.MachineCIDR** range, which is the IP address pool for cluster machines.

- The VNet and its subnets must belong to the same resource group, and the subnets must be configured to use Azure-assigned DHCP IP addresses instead of static IP addresses.

You must provide two subnets within your VNet, one for the control plane machines and one for the compute machines. Because Azure distributes machines in different availability zones within the region that you specify, your cluster will have high availability by default.




NOTE

By default, if you specify availability zones in the **install-config.yaml** file, the installation program distributes the control plane machines and the compute machines across [these availability zones](#) within [a region](#). To ensure high availability for your cluster, select a region with at least three availability zones. If your region contains fewer than three availability zones, the installation program places more than one control plane machine in the available zones.

To ensure that the subnets that you provide are suitable, the installation program confirms the following data:

- All the specified subnets exist.
- There are two private subnets, one for the control plane machines and one for the compute machines.
- The subnet CIDRs belong to the machine CIDR that you specified. Machines are not provisioned in availability zones that you do not provide private subnets for. If required, the installation program creates public load balancers that manage the control plane and worker nodes, and Azure allocates a public IP address to them.




NOTE

If you destroy a cluster that uses an existing VNet, the VNet is not deleted.

7.7.2.1.1. Network security group requirements

The network security groups for the subnets that host the compute and control plane machines require specific access to ensure that the cluster communication is correct. You must create rules to allow access to the required cluster communication ports.



IMPORTANT

The network security group rules must be in place before you install the cluster. If you attempt to install a cluster without the required access, the installation program cannot reach the Azure APIs, and installation fails.

Table 7.14. Required ports

Port	Description	Control plane	Compute
80	Allows HTTP traffic		x
443	Allows HTTPS traffic		x

Port	Description	Control plane	Compute
6443	Allows communication to the control plane machines	x	
22623	Allows internal communication to the machine config server for provisioning machines	x	

1. If you are using Azure Firewall to restrict the internet access, then [you can configure Azure Firewall to allow the Azure APIs](#). A network security group rule is not needed.



IMPORTANT

Currently, there is no supported way to block or restrict the machine config server endpoint. The machine config server must be exposed to the network so that newly-provisioned machines, which have no existing configuration or state, are able to fetch their configuration. In this model, the root of trust is the certificate signing requests (CSR) endpoint, which is where the kubelet sends its certificate signing request for approval to join the cluster. Because of this, machine configs should not be used to distribute sensitive information, such as secrets and certificates.

To ensure that the machine config server endpoints, ports 22623 and 22624, are secured in bare metal scenarios, customers must configure proper network policies.

Because cluster components do not modify the user-provided network security groups, which the Kubernetes controllers update, a pseudo-network security group is created for the Kubernetes controller to modify without impacting the rest of the environment.

Additional resources

- [About the OpenShift SDN network plugin](#)
- [Configuring your firewall](#)

7.7.2.2. Division of permissions

Starting with OpenShift Container Platform 4.3, you do not need all of the permissions that are required for an installation program-provisioned infrastructure cluster to deploy a cluster. This change mimics the division of permissions that you might have at your company: some individuals can create different resources in your clouds than others. For example, you might be able to create application-specific items, like instances, storage, and load balancers, but not networking-related components such as VNets, subnet, or ingress rules.

The Azure credentials that you use when you create your cluster do not need the networking permissions that are required to make VNets and core networking components within the VNet, such as subnets, routing tables, internet gateways, NAT, and VPN. You still need permission to make the application resources that the machines within the cluster require, such as load balancers, security groups, storage accounts, and nodes.

7.7.2.3. Isolation between clusters

Because the cluster is unable to modify network security groups in an existing subnet, there is no way to isolate clusters from each other on the VNet.

7.7.3. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.



IMPORTANT

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

7.7.4. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.



NOTE

You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.



NOTE

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.



NOTE

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```



NOTE

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1 Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

7.7.5. Obtaining the installation program


Before you install OpenShift Container Platform, download the installation file on the host you are using for installation.

Prerequisites

- You have a computer that runs Linux or macOS, with 500 MB of local disk space.


Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select your infrastructure provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.



IMPORTANT

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.



IMPORTANT

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.

4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#) . This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

7.7.6. Creating the installation configuration file

You can customize the OpenShift Container Platform cluster you install on Microsoft Azure.

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have an Azure subscription ID and tenant ID.
- If you are installing the cluster using a service principal, you have its application ID and password.
- If you are installing the cluster using a system-assigned managed identity, you have enabled it on the virtual machine that you will run the installation program from.
- If you are installing the cluster using a user-assigned managed identity, you have met these prerequisites:
 - You have its client ID.
 - You have assigned it to the virtual machine that you will run the installation program from.

Procedure

1. Optional: If you have run the installation program on this computer before, and want to use an alternative service principal or managed identity, go to the `~/.azure/` directory and delete the **osServicePrincipal.json** configuration file.
Deleting this file prevents the installation program from automatically reusing subscription and authentication values from a previous installation.
2. Create the **install-config.yaml** file.
 - a. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install create install-config --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the directory name to store the files that the installation program creates.


When specifying the directory:

- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
 - Use an empty directory. Some installation assets, such as bootstrap X.509 certificates, have short expiration intervals, therefore you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.
- b. At the prompts, provide the configuration details for your cloud:
 - i. Optional: Select an SSH key to use to access your cluster machines.

**NOTE**

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- ii. Select **azure** as the platform to target.
If the installation program cannot locate the **osServicePrincipal.json** configuration file from a previous installation, you are prompted for Azure subscription and authentication values.
- iii. Enter the following Azure parameter values for your subscription:
 - **azure subscription id** Enter the subscription ID to use for the cluster.
 - **azure tenant id** Enter the tenant ID.
- iv. Depending on the Azure identity you are using to deploy the cluster, do one of the following when prompted for the **azure service principal client id**
 - If you are using a service principal, enter its application ID.
 - If you are using a system-assigned managed identity, leave this value blank.
 - If you are using a user-assigned managed identity, specify its client ID.
- v. Depending on the Azure identity you are using to deploy the cluster, do one of the following when prompted for the **azure service principal client secret**
 - If you are using a service principal, enter its password.
 - If you are using a system-assigned managed identity, leave this value blank.
 - If you are using a user-assigned managed identity, leave this value blank.
- vi. Select the region to deploy the cluster to.
- vii. Select the base domain to deploy the cluster to. The base domain corresponds to the Azure DNS Zone that you created for your cluster.
- viii. Enter a descriptive name for your cluster.

**IMPORTANT**

All Azure resources that are available through public endpoints are subject to resource name restrictions, and you cannot create resources that use certain terms. For a list of terms that Azure restricts, see [Resolve reserved resource name errors](#) in the Azure documentation.

3. Modify the **install-config.yaml** file. You can find more information about the available parameters in the "Installation configuration parameters" section.
4. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

If previously not detected, the installation program creates an **osServicePrincipal.json** configuration file and stores this file in the `~/.azure/` directory on your computer. This ensures that the installation program can load the profile when it is creating an OpenShift Container Platform cluster on the target platform.

Additional resources

- [Installation configuration parameters for Azure](#)

7.7.6.1. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 7.15. Minimum resource requirements

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Bootstrap	RHCOS	4	16 GB	100 GB	300
Control plane	RHCOS	4	16 GB	100 GB	300
Compute	RHCOS, RHEL 8.6 and later [3]	2	8 GB	100 GB	300

1. One vCPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio: $(\text{threads per core} \times \text{cores}) \times \text{sockets} = \text{vCPUs}$.
2. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes which require a 10 ms p99 fsync duration. Note that on many cloud platforms, storage size and IOPS scale together, so you might need to over-allocate storage volume to obtain sufficient performance.
3. As with all user-provisioned installations, if you choose to use RHEL compute machines in your cluster, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Use of RHEL 7 compute machines is deprecated and has been removed in OpenShift Container Platform 4.10 and later.




NOTE

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).



IMPORTANT

You are required to use Azure virtual machines that have the **premiumIO** parameter set to **true**.

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

7.7.6.2. Tested instance types for Azure

The following Microsoft Azure instance types have been tested with OpenShift Container Platform.

Example 7.30. Machine types based on 64-bit x86 architecture

- **c4.***
- **c5.***
- **c5a.***
- **i3.***
- **m4.***
- **m5.***
- **m5a.***
- **m6a.***
- **m6i.***
- **r4.***
- **r5.***

- **r5a.***
- **r6i.***
- **t3.***
- **t3a.***

7.7.6.3. Tested instance types for Azure on 64-bit ARM infrastructures

The following Microsoft Azure ARM64 instance types have been tested with OpenShift Container Platform.

Example 7.31. Machine types based on 64-bit ARM architecture

- **c6g.***
- **m6g.***

7.7.6.4. Enabling trusted launch for Azure VMs

You can enable two trusted launch features when installing your cluster on Azure: [secure boot](#) and [virtualized Trusted Platform Modules](#).

See the Azure documentation about [virtual machine sizes](#) to learn what sizes of virtual machines support these features.



IMPORTANT

Trusted launch is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

Prerequisites

- You have created an **install-config.yaml** file.

Procedure

- Use a text editor to edit the **install-config.yaml** file prior to deploying your cluster and add the following stanza:

```
controlPlane: ❶
platform:
  azure:
    settings:
      securityType: TrustedLaunch ❷
```

```

trustedLaunch:
  uefiSettings:
    secureBoot: Enabled 3
    virtualizedTrustedPlatformModule: Enabled 4

```

- 1** Specify **controlPlane.platform.azure** or **compute.platform.azure** to enable trusted launch on only control plane or compute nodes respectively. Specify **platform.azure.defaultMachinePlatform** to enable trusted launch on all nodes.
- 2** Enable trusted launch features.
- 3** Enable secure boot. For more information, see the Azure documentation about [secure boot](#).
- 4** Enable the virtualized Trusted Platform Module. For more information, see the Azure documentation about [virtualized Trusted Platform Modules](#).

7.7.6.5. Enabling confidential VMs

You can enable confidential VMs when installing your cluster. You can enable confidential VMs for compute nodes, control plane nodes, or all nodes.




IMPORTANT

Using confidential VMs is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

You can use confidential VMs with the following VM sizes:

- DCasv5-series
- DCadsv5-series
- ECasv5-series
- ECadsv5-series



IMPORTANT

Confidential VMs are currently not supported on 64-bit ARM architectures.

Prerequisites

- You have created an **install-config.yaml** file.

Procedure

- Use a text editor to edit the **install-config.yaml** file prior to deploying your cluster and add the following stanza:

```
controlPlane: ❶
  platform:
    azure:
      settings:
        securityType: ConfidentialVM ❷
        confidentialVM:
          uefiSettings:
            secureBoot: Enabled ❸
            virtualizedTrustedPlatformModule: Enabled ❹
      osDisk:
        securityProfile:
          securityEncryptionType: VMGuestStateOnly ❺
```

- ❶ Specify **controlPlane.platform.azure** or **compute.platform.azure** to deploy confidential VMs on only control plane or compute nodes respectively. Specify **platform.azure.defaultMachinePlatform** to deploy confidential VMs on all nodes.
- ❷ Enable confidential VMs.
- ❸ Enable secure boot. For more information, see the Azure documentation about [secure boot](#).
- ❹ Enable the virtualized Trusted Platform Module. For more information, see the Azure documentation about [virtualized Trusted Platform Modules](#).
- ❺ Specify **VMGuestStateOnly** to encrypt the VM guest state.

7.7.6.6. Sample customized install-config.yaml file for Azure

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.



IMPORTANT

This sample YAML file is provided for reference only. You must obtain your **install-config.yaml** file by using the installation program and modify it.

```
apiVersion: v1
baseDomain: example.com ❶
controlPlane: ❷
  hyperthreading: Enabled ❸ ❹
  name: master
  platform:
    azure:
      encryptionAtHost: true
      ultraSSDCapability: Enabled
      osDisk:
        diskSizeGB: 1024 ❺
        diskType: Premium_LRS
        diskEncryptionSet:
```

```

    resourceGroup: disk_encryption_set_resource_group
    name: disk_encryption_set_name
    subscriptionId: secondary_subscription_id
  osImage:
    publisher: example_publisher_name
    offer: example_image_offer
    sku: example_offer_sku
    version: example_image_version
    type: Standard_D8s_v3
  replicas: 3
compute: 6
- hyperthreading: Enabled 7
  name: worker
  platform:
    azure:
      ultraSSDCapability: Enabled
      type: Standard_D2s_v3
      encryptionAtHost: true
    osDisk:
      diskSizeGB: 512 8
      diskType: Standard_LRS
      diskEncryptionSet:
        resourceGroup: disk_encryption_set_resource_group
        name: disk_encryption_set_name
        subscriptionId: secondary_subscription_id
    osImage:
      publisher: example_publisher_name
      offer: example_image_offer
      sku: example_offer_sku
      version: example_image_version
  zones: 9
  - "1"
  - "2"
  - "3"
  replicas: 5
metadata:
  name: test-cluster 10
networking:
  clusterNetwork:
    - cidr: 10.128.0.0/14
    hostPrefix: 23
  machineNetwork:
    - cidr: 10.0.0.0/16
  networkType: OVNKubernetes 11
  serviceNetwork:
    - 172.30.0.0/16
platform:
  azure:
    defaultMachinePlatform:
      osImage: 12
      publisher: example_publisher_name
      offer: example_image_offer
      sku: example_offer_sku
      version: example_image_version
      ultraSSDCapability: Enabled

```

```

baseDomainResourceGroupName: resource_group 13
region: centralus 14
resourceGroupName: existing_resource_group 15
networkResourceGroupName: vnet_resource_group 16
virtualNetwork: vnet 17
controlPlaneSubnet: control_plane_subnet 18
computeSubnet: compute_subnet 19
outboundType: Loadbalancer
cloudName: AzurePublicCloud
pullSecret: '{"auths": ...}' 20
fips: false 21
sshKey: ssh-ed25519 AAAA... 22

```

- 1 10 14 20 Required. The installation program prompts you for this value.
- 2 6 If you do not provide these parameters and values, the installation program provides the default value.
- 3 7 The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.
- 4 Whether to enable or disable simultaneous multithreading, or **hyperthreading**. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores. You can disable it by setting the parameter value to **Disabled**. If you disable simultaneous multithreading in some cluster machines, you must disable it in all cluster machines.



IMPORTANT

If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance. Use larger virtual machine types, such as **Standard_D8s_v3**, for your machines if you disable simultaneous multithreading.

- 5 8 You can specify the size of the disk to use in GB. Minimum recommendation for control plane nodes is 1024 GB.
- 9 Specify a list of zones to deploy your machines to. For high availability, specify at least two zones.
- 11 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 12 Optional: A custom Red Hat Enterprise Linux CoreOS (RHCOS) image that should be used to boot control plane and compute machines. The **publisher**, **offer**, **sku**, and **version** parameters under **platform.azure.defaultMachinePlatform.osImage** apply to both control plane and compute machines. If the parameters under **controlPlane.platform.azure.osImage** or **compute.platform.azure.osImage** are set, they override the **platform.azure.defaultMachinePlatform.osImage** parameters.
- 13 Specify the name of the resource group that contains the DNS zone for your base domain.
- 15 Specify the name of an already existing resource group to install your cluster to. If undefined, a new resource group is created for the cluster.

- 16 If you use an existing VNet, specify the name of the resource group that contains it.
- 17 If you use an existing VNet, specify its name.
- 18 If you use an existing VNet, specify the name of the subnet to host the control plane machines.
- 19 If you use an existing VNet, specify the name of the subnet to host the compute machines.
- 21 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.



IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 22 You can optionally provide the **sshKey** value that you use to access the machines in your cluster.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

7.7.6.7. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.



NOTE

The installation program does not support the proxy **readinessEndpoints** field.

**NOTE**

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.

**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

Additional resources

- For more details about Accelerated Networking, see [Accelerated Networking for Microsoft Azure VMs](#).

7.7.7. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.

**IMPORTANT**

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**.
To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.



NOTE

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your **PATH**.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

7.7.8. Alternatives to storing administrator-level secrets in the kube-system project

By default, administrator secrets are stored in the **kube-system** project. If you configured the **credentialsMode** parameter in the **install-config.yaml** file to **Manual**, you must use one of the following alternatives:

- To manage long-term cloud credentials manually, follow the procedure in [Manually creating long-term credentials](#).
- To implement short-term credentials that are managed outside the cluster for individual components, follow the procedures in [Configuring an Azure cluster to use short-term credentials](#).

7.7.8.1. Manually creating long-term credentials

The Cloud Credential Operator (CCO) can be put into manual mode prior to installation in environments where the cloud identity and access management (IAM) APIs are not reachable, or the administrator prefers not to store an administrator-level credential secret in the cluster **kube-system** namespace.

Procedure

1. If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```
apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...
```

2. If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

3. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

4. Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included \1
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \2
  --to=<path_to_directory_for_credentials_requests> \3
```

- 1 The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2 Specify the location of the **install-config.yaml** file.
- 3 Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

This command creates a YAML file for each **CredentialsRequest** object.

Sample **CredentialsRequest** object

```
apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: AzureProviderSpec
    roleBindings:
      - role: Contributor
    ...
```

5. Create YAML files for secrets in the **openshift-install** manifests directory that you generated previously. The secrets must be stored using the namespace and secret name defined in the **spec.secretRef** for each **CredentialsRequest** object.

Sample **CredentialsRequest** object with secrets

```
apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: AzureProviderSpec
    roleBindings:
      - role: Contributor
```

```

...
secretRef:
  name: <component_secret>
  namespace: <component_namespace>
...


```

Sample Secret object

```

apiVersion: v1
kind: Secret
metadata:
  name: <component_secret>
  namespace: <component_namespace>
data:
  azure_subscription_id: <base64_encoded_azure_subscription_id>
  azure_client_id: <base64_encoded_azure_client_id>
  azure_client_secret: <base64_encoded_azure_client_secret>
  azure_tenant_id: <base64_encoded_azure_tenant_id>
  azure_resource_prefix: <base64_encoded_azure_resource_prefix>
  azure_resourcegroup: <base64_encoded_azure_resourcegroup>
  azure_region: <base64_encoded_azure_region>

```



IMPORTANT

Before upgrading a cluster that uses manually maintained credentials, you must ensure that the CCO is in an upgradeable state.

7.7.8.2. Configuring an Azure cluster to use short-term credentials

To install a cluster that uses Microsoft Entra Workload ID, you must configure the Cloud Credential Operator utility and create the required Azure resources for your cluster.

7.7.8.2.1. Configuring the Cloud Credential Operator utility

To create and manage cloud credentials from outside of the cluster when the Cloud Credential Operator (CCO) is operating in manual mode, extract and prepare the CCO utility (**ccoctl**) binary.



NOTE

The **ccoctl** utility is a Linux binary that must run in a Linux environment.

Prerequisites

- You have access to an OpenShift Container Platform account with cluster administrator access.
- You have installed the OpenShift CLI (**oc**).
- You have created a global Microsoft Azure account for the **ccoctl** utility to use with the following permissions:

Example 7.32. Required Azure permissions

- Microsoft.Resources/subscriptions/resourceGroups/read
- Microsoft.Resources/subscriptions/resourceGroups/write

- Microsoft.Resources/subscriptions/resourceGroups/delete
- Microsoft.Authorization/roleAssignments/read
- Microsoft.Authorization/roleAssignments/delete
- Microsoft.Authorization/roleAssignments/write
- Microsoft.Authorization/roleDefinitions/read
- Microsoft.Authorization/roleDefinitions/write
- Microsoft.Authorization/roleDefinitions/delete
- Microsoft.Storage/storageAccounts/listkeys/action
- Microsoft.Storage/storageAccounts/delete
- Microsoft.Storage/storageAccounts/read
- Microsoft.Storage/storageAccounts/write
- Microsoft.Storage/storageAccounts/blobServices/containers/write
- Microsoft.Storage/storageAccounts/blobServices/containers/delete
- Microsoft.Storage/storageAccounts/blobServices/containers/read
- Microsoft.ManagedIdentity/userAssignedIdentities/delete
- Microsoft.ManagedIdentity/userAssignedIdentities/read
- Microsoft.ManagedIdentity/userAssignedIdentities/write
- Microsoft.ManagedIdentity/userAssignedIdentities/federatedIdentityCredentials/read
- Microsoft.ManagedIdentity/userAssignedIdentities/federatedIdentityCredentials/write
- Microsoft.ManagedIdentity/userAssignedIdentities/federatedIdentityCredentials/delete
- Microsoft.Storage/register/action
- Microsoft.ManagedIdentity/register/action

Procedure

1. Set a variable for the OpenShift Container Platform release image by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

2. Obtain the CCO container image from the OpenShift Container Platform release image by running the following command:

```
$ CCO_IMAGE=$(oc adm release info --image-for='cloud-credential-operator'  
$RELEASE_IMAGE -a ~/.pull-secret)
```

**NOTE**

Ensure that the architecture of the **\$RELEASE_IMAGE** matches the architecture of the environment in which you will use the **ccoctl** tool.

3. Extract the **ccoctl** binary from the CCO container image within the OpenShift Container Platform release image by running the following command:

```
$ oc image extract $CCO_IMAGE \
  --file="/usr/bin/ccoctl.<rhel_version>" \
  -a ~/.pull-secret
```

- 1 For **<rhel_version>**, specify the value that corresponds to the version of Red Hat Enterprise Linux (RHEL) that the host uses. If no value is specified, **ccoctl.rhel8** is used by default. The following values are valid:

- **rhel8**: Specify this value for hosts that use RHEL 8.
- **rhel9**: Specify this value for hosts that use RHEL 9.

4. Change the permissions to make **ccoctl** executable by running the following command:

```
$ chmod 775 ccoctl.<rhel_version>
```

Verification

- To verify that **ccoctl** is ready to use, display the help file by running the following command:

```
$ ccoctl --help
```

Output of ccoctl --help

```
OpenShift credentials provisioning tool
```

```
Usage:
```

```
ccoctl [command]
```

```
Available Commands:
```

```
aws      Manage credentials objects for AWS cloud
azure    Manage credentials objects for Azure
gcp      Manage credentials objects for Google cloud
help     Help about any command
ibmcloud Manage credentials objects for IBM Cloud
nutanix  Manage credentials objects for Nutanix
```

```
Flags:
```

```
-h, --help  help for ccoctl
```

```
Use "ccoctl [command] --help" for more information about a command.
```

7.7.8.2.2. Creating Azure resources with the Cloud Credential Operator utility

You can use the **ccoctl azure create-all** command to automate the creation of Azure resources.



NOTE

By default, **ccoctl** creates objects in the directory in which the commands are run. To create the objects in a different directory, use the **--output-dir** flag. This procedure uses **<path_to_ccoctl_output_dir>** to refer to this directory.

Prerequisites

You must have:

- Extracted and prepared the **ccoctl** binary.
- Access to your Microsoft Azure account by using the Azure CLI.

Procedure

1. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

2. Extract the list of **CredentialsRequest** objects from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included \1 \
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \2 \
  --to=<path_to_directory_for_credentials_requests> \3
```

- 1 The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2 Specify the location of the **install-config.yaml** file.
- 3 Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.



NOTE

This command might take a few moments to run.

3. To enable the **ccoctl** utility to detect your Azure credentials automatically, log in to the Azure CLI by running the following command:

```
$ az login
```

4. Use the **ccoctl** tool to process all **CredentialsRequest** objects by running the following command:

■


```
$ ccoctl azure create-all \
  --name=<azure_infra_name> \ 1
  --output-dir=<ccoctl_output_dir> \ 2
  --region=<azure_region> \ 3
  --subscription-id=<azure_subscription_id> \ 4
  --credentials-requests-dir=<path_to_credentials_requests_directory> \ 5
  --dnszone-resource-group-name=<azure_dns_zone_resource_group_name> \ 6
  --tenant-id=<azure_tenant_id> \ 7
```

- 1 Specify the user-defined name for all created Azure resources used for tracking.
- 2 Optional: Specify the directory in which you want the **ccoctl** utility to create objects. By default, the utility creates objects in the directory in which the commands are run.
- 3 Specify the Azure region in which cloud resources will be created.
- 4 Specify the Azure subscription ID to use.
- 5 Specify the directory containing the files for the component **CredentialsRequest** objects.
- 6 Specify the name of the resource group containing the cluster's base domain Azure DNS zone.
- 7 Specify the Azure tenant ID to use.



NOTE

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

To see additional optional parameters and explanations of how to use them, run the **azure create-all --help** command.

Verification

- To verify that the OpenShift Container Platform secrets are created, list the files in the **<path_to_ccoctl_output_dir>/manifests** directory:

```
$ ls <path_to_ccoctl_output_dir>/manifests
```

Example output

```
azure-ad-pod-identity-webhook-config.yaml
cluster-authentication-02-config.yaml
openshift-cloud-controller-manager-azure-cloud-credentials-credentials.yaml
openshift-cloud-network-config-controller-cloud-credentials-credentials.yaml
openshift-cluster-api-capz-manager-bootstrap-credentials-credentials.yaml
openshift-cluster-csi-drivers-azure-disk-credentials-credentials.yaml
openshift-cluster-csi-drivers-azure-file-credentials-credentials.yaml
openshift-image-registry-installer-cloud-credentials-credentials.yaml
openshift-ingress-operator-cloud-credentials-credentials.yaml
openshift-machine-api-azure-cloud-credentials-credentials.yaml
```

You can verify that the Microsoft Entra ID service accounts are created by querying Azure. For more information, refer to Azure documentation on listing Entra ID service accounts.

7.7.8.2.3. Incorporating the Cloud Credential Operator utility manifests

To implement short-term security credentials managed outside the cluster for individual components, you must move the manifest files that the Cloud Credential Operator utility (**ccoctl**) created to the correct directories for the installation program.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have configured the Cloud Credential Operator utility (**ccoctl**).
- You have created the cloud provider resources that are required for your cluster with the **ccoctl** utility.

Procedure

1. If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```
apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...
```

2. If you used the **ccoctl** utility to create a new Azure resource group instead of using an existing resource group, modify the **resourceGroupName** parameter in the **install-config.yaml** as shown:

Sample configuration file snippet

```
apiVersion: v1
baseDomain: example.com
# ...
platform:
  azure:
    resourceGroupName: <azure_infra_name> 1
# ...
```

- 1** This value must match the user-defined name for Azure resources that was specified with the **--name** argument of the **ccoctl azure create-all** command.

3. If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

- Copy the manifests that the **ccoctl** utility generated to the **manifests** directory that the installation program created by running the following command:

```
$ cp /<path_to_ccoctl_output_dir>/manifests/* ./manifests/
```

- Copy the private key that the **ccoctl** utility generated in the **tls** directory to the installation directory by running the following command:

```
$ cp -a /<path_to_ccoctl_output_dir>/tls .
```

7.7.9. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have an Azure subscription ID and tenant ID.

Procedure

- Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ ❶  
--log-level=info ❷
```

- For **<installation_directory>**, specify the location of your customized **./install-config.yaml** file.
- To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to **<installation_directory>/openshift_install.log**.

**IMPORTANT**

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```

**IMPORTANT**

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

Additional resources

- See [Accessing the web console](#) for more details about accessing and understanding the OpenShift Container Platform web console.

7.7.10. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service

7.7.11. Next steps

- [Customize your cluster](#).

- If necessary, you can [opt out of remote health reporting](#) .

7.8. INSTALLING A PRIVATE CLUSTER ON AZURE

In OpenShift Container Platform version 4.16, you can install a private cluster into an existing Azure Virtual Network (VNet) on Microsoft Azure. The installation program provisions the rest of the required infrastructure, which you can further customize. To customize the installation, you modify parameters in the **install-config.yaml** file before you install the cluster.

7.8.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [configured an Azure account](#) to host the cluster and determined the tested and validated region to deploy the cluster to.
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.
- If you use customer-managed encryption keys, you [prepared your Azure environment for encryption](#).

7.8.2. Private clusters

You can deploy a private OpenShift Container Platform cluster that does not expose external endpoints. Private clusters are accessible from only an internal network and are not visible to the internet.

By default, OpenShift Container Platform is provisioned to use publicly-accessible DNS and endpoints. A private cluster sets the DNS, Ingress Controller, and API server to private when you deploy your cluster. This means that the cluster resources are only accessible from your internal network and are not visible to the internet.



IMPORTANT

If the cluster has any public subnets, load balancer services created by administrators might be publicly accessible. To ensure cluster security, verify that these services are explicitly annotated as private.

To deploy a private cluster, you must:

- Use existing networking that meets your requirements. Your cluster resources might be shared between other clusters on the network.
- Deploy from a machine that has access to:
 - The API services for the cloud to which you provision.
 - The hosts on the network that you provision.
 - The internet to obtain installation media.

You can use any machine that meets these access requirements and follows your company's guidelines. For example, this machine can be a bastion host on your cloud network or a machine that has access to the network through a VPN.

7.8.2.1. Private clusters in Azure

To create a private cluster on Microsoft Azure, you must provide an existing private VNet and subnets to host the cluster. The installation program must also be able to resolve the DNS records that the cluster requires. The installation program configures the Ingress Operator and API server for only internal traffic.

Depending how your network connects to the private VNET, you might need to use a DNS forwarder to resolve the cluster's private DNS records. The cluster's machines use **168.63.129.16** internally for DNS resolution. For more information, see [What is Azure Private DNS?](#) and [What is IP address 168.63.129.16?](#) in the Azure documentation.

The cluster still requires access to internet to access the Azure APIs.

The following items are not required or created when you install a private cluster:

- A **BaseDomainResourceGroup**, since the cluster does not create public records
- Public IP addresses
- Public DNS records
- Public endpoints

The cluster is configured so that the Operators do not create public records for the cluster and all cluster machines are placed in the private subnets that you specify.

7.8.2.1.1. Limitations

Private clusters on Azure are subject to only the limitations that are associated with the use of an existing VNet.

7.8.2.2. User-defined outbound routing

In OpenShift Container Platform, you can choose your own outbound routing for a cluster to connect to the internet. This allows you to skip the creation of public IP addresses and the public load balancer.

You can configure user-defined routing by modifying parameters in the **install-config.yaml** file before installing your cluster. A pre-existing VNet is required to use outbound routing when installing a cluster; the installation program is not responsible for configuring this.

When configuring a cluster to use user-defined routing, the installation program does not create the following resources:

- Outbound rules for access to the internet.
- Public IPs for the public load balancer.
- Kubernetes Service object to add the cluster machines to the public load balancer for outbound requests.

You must ensure the following items are available before setting user-defined routing:

- Egress to the internet is possible to pull container images, unless using an OpenShift image registry mirror.
- The cluster can access Azure APIs.
- Various allowlist endpoints are configured. You can reference these endpoints in the *Configuring your firewall* section.

There are several pre-existing networking setups that are supported for internet access using user-defined routing.

Private cluster with network address translation

You can use [Azure VNET network address translation \(NAT\)](#) to provide outbound internet access for the subnets in your cluster. You can reference [Create a NAT gateway using Azure CLI](#) in the Azure documentation for configuration instructions.

When using a VNet setup with Azure NAT and user-defined routing configured, you can create a private cluster with no public endpoints.

Private cluster with Azure Firewall

You can use Azure Firewall to provide outbound routing for the VNet used to install the cluster. You can learn more about [providing user-defined routing with Azure Firewall](#) in the Azure documentation.

When using a VNet setup with Azure Firewall and user-defined routing configured, you can create a private cluster with no public endpoints.

Private cluster with a proxy configuration

You can use a proxy with user-defined routing to allow egress to the internet. You must ensure that cluster Operators do not access Azure APIs using a proxy; Operators must have access to Azure APIs outside of the proxy.

When using the default route table for subnets, with **0.0.0.0/0** populated automatically by Azure, all Azure API requests are routed over Azure's internal network even though the IP addresses are public. As long as the Network Security Group rules allow egress to Azure API endpoints, proxies with user-defined routing configured allow you to create private clusters with no public endpoints.

Private cluster with no internet access

You can install a private network that restricts all access to the internet, except the Azure API. This is accomplished by mirroring the release image registry locally. Your cluster must have access to the following:

- An OpenShift image registry mirror that allows for pulling container images
- Access to Azure APIs

With these requirements available, you can use user-defined routing to create private clusters with no public endpoints.

7.8.3. About reusing a VNet for your OpenShift Container Platform cluster

In OpenShift Container Platform 4.16, you can deploy a cluster into an existing Azure Virtual Network (VNet) in Microsoft Azure. If you do, you must also use existing subnets within the VNet and routing rules.

By deploying OpenShift Container Platform into an existing Azure VNet, you might be able to avoid service limit constraints in new accounts or more easily abide by the operational constraints that your company's guidelines set. This is a good option to use if you cannot obtain the infrastructure creation

permissions that are required to create the VNet.

7.8.3.1. Requirements for using your VNet

When you deploy a cluster by using an existing VNet, you must perform additional network configuration before you install the cluster. In installer-provisioned infrastructure clusters, the installer usually creates the following components, but it does not create them when you install into an existing VNet:

- Subnets
- Route tables
- VNets
- Network Security Groups



NOTE

The installation program requires that you use the cloud-provided DNS server. Using a custom DNS server is not supported and causes the installation to fail.

If you use a custom VNet, you must correctly configure it and its subnets for the installation program and the cluster to use. The installation program cannot subdivide network ranges for the cluster to use, set route tables for the subnets, or set VNet options like DHCP, so you must do so before you install the cluster.

The cluster must be able to access the resource group that contains the existing VNet and subnets. While all of the resources that the cluster creates are placed in a separate resource group that it creates, some network resources are used from a separate group. Some cluster Operators must be able to access resources in both resource groups. For example, the Machine API controller attaches NICs for the virtual machines that it creates to subnets from the networking resource group.

Your VNet must meet the following characteristics:

- The VNet's CIDR block must contain the **Networking.MachineCIDR** range, which is the IP address pool for cluster machines.
- The VNet and its subnets must belong to the same resource group, and the subnets must be configured to use Azure-assigned DHCP IP addresses instead of static IP addresses.

You must provide two subnets within your VNet, one for the control plane machines and one for the compute machines. Because Azure distributes machines in different availability zones within the region that you specify, your cluster will have high availability by default.



NOTE

By default, if you specify availability zones in the **install-config.yaml** file, the installation program distributes the control plane machines and the compute machines across [these availability zones](#) within [a region](#). To ensure high availability for your cluster, select a region with at least three availability zones. If your region contains fewer than three availability zones, the installation program places more than one control plane machine in the available zones.

To ensure that the subnets that you provide are suitable, the installation program confirms the following data:

- All the specified subnets exist.
- There are two private subnets, one for the control plane machines and one for the compute machines.
- The subnet CIDRs belong to the machine CIDR that you specified. Machines are not provisioned in availability zones that you do not provide private subnets for.

**NOTE**

If you destroy a cluster that uses an existing VNet, the VNet is not deleted.

7.8.3.1.1. Network security group requirements

The network security groups for the subnets that host the compute and control plane machines require specific access to ensure that the cluster communication is correct. You must create rules to allow access to the required cluster communication ports.

**IMPORTANT**

The network security group rules must be in place before you install the cluster. If you attempt to install a cluster without the required access, the installation program cannot reach the Azure APIs, and installation fails.

Table 7.16. Required ports

Port	Description	Control plane	Compute
80	Allows HTTP traffic		x
443	Allows HTTPS traffic		x
6443	Allows communication to the control plane machines	x	
22623	Allows internal communication to the machine config server for provisioning machines	x	

1. If you are using Azure Firewall to restrict the internet access, then [you can configure Azure Firewall to allow the Azure APIs](#). A network security group rule is not needed.

**IMPORTANT**

Currently, there is no supported way to block or restrict the machine config server endpoint. The machine config server must be exposed to the network so that newly-provisioned machines, which have no existing configuration or state, are able to fetch their configuration. In this model, the root of trust is the certificate signing requests (CSR) endpoint, which is where the kubelet sends its certificate signing request for approval to join the cluster. Because of this, machine configs should not be used to distribute sensitive information, such as secrets and certificates.

To ensure that the machine config server endpoints, ports 22623 and 22624, are secured in bare metal scenarios, customers must configure proper network policies.

Because cluster components do not modify the user-provided network security groups, which the Kubernetes controllers update, a pseudo-network security group is created for the Kubernetes controller to modify without impacting the rest of the environment.

Additional resources

- [About the OpenShift SDN network plugin](#)
- [Configuring your firewall](#)

7.8.3.2. Division of permissions

Starting with OpenShift Container Platform 4.3, you do not need all of the permissions that are required for an installation program-provisioned infrastructure cluster to deploy a cluster. This change mimics the division of permissions that you might have at your company: some individuals can create different resources in your clouds than others. For example, you might be able to create application-specific items, like instances, storage, and load balancers, but not networking-related components such as VNets, subnet, or ingress rules.

The Azure credentials that you use when you create your cluster do not need the networking permissions that are required to make VNets and core networking components within the VNet, such as subnets, routing tables, internet gateways, NAT, and VPN. You still need permission to make the application resources that the machines within the cluster require, such as load balancers, security groups, storage accounts, and nodes.

7.8.3.3. Isolation between clusters

Because the cluster is unable to modify network security groups in an existing subnet, there is no way to isolate clusters from each other on the VNet.

7.8.4. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.



IMPORTANT


If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

7.8.5. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.



NOTE

You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.



NOTE

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.

**NOTE**

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```

**NOTE**

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

7.8.6. Obtaining the installation program


Before you install OpenShift Container Platform, download the installation file on the host you are using for installation.

Prerequisites

- You have a computer that runs Linux or macOS, with 500 MB of local disk space.


Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select your infrastructure provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.



IMPORTANT

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.



IMPORTANT

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.

4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#) . This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

7.8.7. Manually creating the installation configuration file

Installing the cluster requires that you manually create the installation configuration file.


Prerequisites

- You have an SSH public key on your local machine to provide to the installation program. The key will be used for SSH authentication onto your cluster nodes for debugging and disaster recovery.
- You have obtained the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create an installation directory to store your required installation assets in:

```
$ mkdir <installation_directory>
```

**IMPORTANT**


You must create a directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

2. Customize the sample **install-config.yaml** file template that is provided and save it in the **<installation_directory>**.

**NOTE**

You must name this configuration file **install-config.yaml**.

3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.

**IMPORTANT**

The **install-config.yaml** file is consumed during the next step of the installation process. You must back it up now.

Additional resources

- [Installation configuration parameters for Azure](#)

7.8.7.1. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 7.17. Minimum resource requirements

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Bootstrap	RHCOS	4	16 GB	100 GB	300
Control plane	RHCOS	4	16 GB	100 GB	300
Compute	RHCOS, RHEL 8.6 and later [3]	2	8 GB	100 GB	300

1. One vCPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio: (threads per core × cores) × sockets = vCPUs.
2. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes which require a 10 ms p99 fsync duration. Note that on many cloud platforms, storage size and IOPS scale together, so

you might need to over-allocate storage volume to obtain sufficient performance.

3. As with all user-provisioned installations, if you choose to use RHEL compute machines in your cluster, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Use of RHEL 7 compute machines is deprecated and has been removed in OpenShift Container Platform 4.10 and later.




NOTE

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).



IMPORTANT

You are required to use Azure virtual machines that have the **premiumIO** parameter set to **true**.

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

7.8.7.2. Tested instance types for Azure

The following Microsoft Azure instance types have been tested with OpenShift Container Platform.

Example 7.33. Machine types based on 64-bit x86 architecture

- **c4.***
- **c5.***
- **c5a.***
- **i3.***
- **m4.***
- **m5.***
- **m5a.***

- **m6a.***
- **m6i.***
- **r4.***
- **r5.***
- **r5a.***
- **r6i.***
- **t3.***
- **t3a.***

7.8.7.3. Tested instance types for Azure on 64-bit ARM infrastructures

The following Microsoft Azure ARM64 instance types have been tested with OpenShift Container Platform.

Example 7.34. Machine types based on 64-bit ARM architecture

- **c6g.***
- **m6g.***

7.8.7.4. Enabling trusted launch for Azure VMs

You can enable two trusted launch features when installing your cluster on Azure: [secure boot](#) and [virtualized Trusted Platform Modules](#).

See the Azure documentation about [virtual machine sizes](#) to learn what sizes of virtual machines support these features.



IMPORTANT

Trusted launch is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

Prerequisites

- You have created an **install-config.yaml** file.

Procedure

- Use a text editor to edit the **install-config.yaml** file prior to deploying your cluster and add the following stanza:

```
controlPlane: ❶
  platform:
    azure:
      settings:
        securityType: TrustedLaunch ❷
      trustedLaunch:
        uefiSettings:
          secureBoot: Enabled ❸
          virtualizedTrustedPlatformModule: Enabled ❹
```

- ❶ Specify **controlPlane.platform.azure** or **compute.platform.azure** to enable trusted launch on only control plane or compute nodes respectively. Specify **platform.azure.defaultMachinePlatform** to enable trusted launch on all nodes.
- ❷ Enable trusted launch features.
- ❸ Enable secure boot. For more information, see the Azure documentation about [secure boot](#).
- ❹ Enable the virtualized Trusted Platform Module. For more information, see the Azure documentation about [virtualized Trusted Platform Modules](#).

7.8.7.5. Enabling confidential VMs

You can enable confidential VMs when installing your cluster. You can enable confidential VMs for compute nodes, control plane nodes, or all nodes.




IMPORTANT

Using confidential VMs is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

You can use confidential VMs with the following VM sizes:

- DCasv5-series
- DCadsv5-series
- ECasv5-series
- ECadsv5-series



IMPORTANT

Confidential VMs are currently not supported on 64-bit ARM architectures.

Prerequisites

- You have created an **install-config.yaml** file.

Procedure

- Use a text editor to edit the **install-config.yaml** file prior to deploying your cluster and add the following stanza:

```
controlPlane: ❶
  platform:
    azure:
      settings:
        securityType: ConfidentialVM ❷
        confidentialVM:
          uefiSettings:
            secureBoot: Enabled ❸
            virtualizedTrustedPlatformModule: Enabled ❹
      osDisk:
        securityProfile:
          securityEncryptionType: VMGuestStateOnly ❺
```

- Specify **controlPlane.platform.azure** or **compute.platform.azure** to deploy confidential VMs on only control plane or compute nodes respectively. Specify **platform.azure.defaultMachinePlatform** to deploy confidential VMs on all nodes.
- Enable confidential VMs.
- Enable secure boot. For more information, see the Azure documentation about [secure boot](#).
- Enable the virtualized Trusted Platform Module. For more information, see the Azure documentation about [virtualized Trusted Platform Modules](#).
- Specify **VMGuestStateOnly** to encrypt the VM guest state.

7.8.7.6. Sample customized install-config.yaml file for Azure

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.



IMPORTANT

This sample YAML file is provided for reference only. You must obtain your **install-config.yaml** file by using the installation program and modify it.

```
apiVersion: v1
baseDomain: example.com ❶
controlPlane: ❷
  hyperthreading: Enabled ❸ ❹
  name: master
  platform:
    azure:
```

```

encryptionAtHost: true
ultraSSDCapability: Enabled
osDisk:
  diskSizeGB: 1024 5
  diskType: Premium_LRS
  diskEncryptionSet:
    resourceGroup: disk_encryption_set_resource_group
    name: disk_encryption_set_name
    subscriptionId: secondary_subscription_id
osImage:
  publisher: example_publisher_name
  offer: example_image_offer
  sku: example_offer_sku
  version: example_image_version
  type: Standard_D8s_v3
replicas: 3
compute: 6
- hyperthreading: Enabled 7
  name: worker
  platform:
    azure:
      ultraSSDCapability: Enabled
      type: Standard_D2s_v3
      encryptionAtHost: true
      osDisk:
        diskSizeGB: 512 8
        diskType: Standard_LRS
        diskEncryptionSet:
          resourceGroup: disk_encryption_set_resource_group
          name: disk_encryption_set_name
          subscriptionId: secondary_subscription_id
      osImage:
        publisher: example_publisher_name
        offer: example_image_offer
        sku: example_offer_sku
        version: example_image_version
    zones: 9
      - "1"
      - "2"
      - "3"
    replicas: 5
  metadata:
    name: test-cluster 10
  networking:
    clusterNetwork:
      - cidr: 10.128.0.0/14
      hostPrefix: 23
    machineNetwork:
      - cidr: 10.0.0.0/16
    networkType: OVNKubernetes 11
    serviceNetwork:
      - 172.30.0.0/16
  platform:
    azure:
      defaultMachinePlatform:

```

```

osImage: 12
  publisher: example_publisher_name
  offer: example_image_offer
  sku: example_offer_sku
  version: example_image_version
  ultraSSDCapability: Enabled
baseDomainResourceGroupName: resource_group 13
region: centralus 14
resourceGroupName: existing_resource_group 15
networkResourceGroupName: vnet_resource_group 16
virtualNetwork: vnet 17
controlPlaneSubnet: control_plane_subnet 18
computeSubnet: compute_subnet 19
outboundType: UserDefinedRouting 20
cloudName: AzurePublicCloud
pullSecret: '{"auths": ...}' 21
fips: false 22
sshKey: ssh-ed25519 AAAA... 23
publish: Internal 24

```

- 1 10 14 21 Required. The installation program prompts you for this value.
- 2 6 If you do not provide these parameters and values, the installation program provides the default value.
- 3 7 The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.
- 4 Whether to enable or disable simultaneous multithreading, or **hyperthreading**. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores. You can disable it by setting the parameter value to **Disabled**. If you disable simultaneous multithreading in some cluster machines, you must disable it in all cluster machines.




IMPORTANT

If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance. Use larger virtual machine types, such as **Standard_D8s_v3**, for your machines if you disable simultaneous multithreading.

- 5 8 You can specify the size of the disk to use in GB. Minimum recommendation for control plane nodes is 1024 GB.
- 9 Specify a list of zones to deploy your machines to. For high availability, specify at least two zones.
- 11 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 12 Optional: A custom Red Hat Enterprise Linux CoreOS (RHCOS) image that should be used to boot control plane and compute machines. The **publisher**, **offer**, **sku**, and **version** parameters under **platform.azure.defaultMachinePlatform.osImage** apply to both control plane and compute machines. If the parameters under **controlPlane.platform.azure.osImage** or

`compute.platform.azure.osImage` are set, they override the `platform.azure.defaultMachinePlatform.osImage` parameters.

- 13 Specify the name of the resource group that contains the DNS zone for your base domain.
- 15 Specify the name of an already existing resource group to install your cluster to. If undefined, a new resource group is created for the cluster.
- 16 If you use an existing VNet, specify the name of the resource group that contains it.
- 17 If you use an existing VNet, specify its name.
- 18 If you use an existing VNet, specify the name of the subnet to host the control plane machines.
- 19 If you use an existing VNet, specify the name of the subnet to host the compute machines.
- 20 You can customize your own outbound routing. Configuring user-defined routing prevents exposing external endpoints in your cluster. User-defined routing for egress requires deploying your cluster to an existing VNet.
- 22 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.



IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 23 You can optionally provide the `sshKey` value that you use to access the machines in your cluster.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your `ssh-agent` process uses.

- 24 How to publish the user-facing endpoints of your cluster. Set `publish` to `Internal` to deploy a private cluster, which cannot be accessed from the internet. The default value is `External`.

7.8.7.7. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the `install-config.yaml` file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.


For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the

**NOTE**

The installation program does not support the proxy **readinessEndpoints** field.

**NOTE**

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.

**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

Additional resources

- For more details about Accelerated Networking, see [Accelerated Networking for Microsoft Azure VMs](#).

7.8.8. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.

**IMPORTANT**

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**. To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**. To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.

**NOTE**

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your PATH.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

7.8.9. Alternatives to storing administrator-level secrets in the kube-system project

By default, administrator secrets are stored in the **kube-system** project. If you configured the **credentialsMode** parameter in the **install-config.yaml** file to **Manual**, you must use one of the following alternatives:

- To manage long-term cloud credentials manually, follow the procedure in [Manually creating long-term credentials](#).
- To implement short-term credentials that are managed outside the cluster for individual components, follow the procedures in [Configuring an Azure cluster to use short-term credentials](#).

7.8.9.1. Manually creating long-term credentials

The Cloud Credential Operator (CCO) can be put into manual mode prior to installation in environments where the cloud identity and access management (IAM) APIs are not reachable, or the administrator prefers not to store an administrator-level credential secret in the cluster **kube-system** namespace.

Procedure

1. If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```
apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...
```

2. If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

- Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

- Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included 1 \
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml 2 \
  --to=<path_to_directory_for_credentials_requests> 3
```

- The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- Specify the location of the **install-config.yaml** file.
- Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

This command creates a YAML file for each **CredentialsRequest** object.

Sample **CredentialsRequest** object

```
apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: AzureProviderSpec
    roleBindings:
      - role: Contributor
    ...
```

- Create YAML files for secrets in the **openshift-install** manifests directory that you generated previously. The secrets must be stored using the namespace and secret name defined in the **spec.secretRef** for each **CredentialsRequest** object.

Sample **CredentialsRequest** object with secrets

```
apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
```

```

namespace: openshift-cloud-credential-operator
...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: AzureProviderSpec
    roleBindings:
      - role: Contributor
    ...
  secretRef:
    name: <component_secret>
    namespace: <component_namespace>
  ...

```

Sample Secret object

```

apiVersion: v1
kind: Secret
metadata:
  name: <component_secret>
  namespace: <component_namespace>
data:
  azure_subscription_id: <base64_encoded_azure_subscription_id>
  azure_client_id: <base64_encoded_azure_client_id>
  azure_client_secret: <base64_encoded_azure_client_secret>
  azure_tenant_id: <base64_encoded_azure_tenant_id>
  azure_resource_prefix: <base64_encoded_azure_resource_prefix>
  azure_resourcegroup: <base64_encoded_azure_resourcegroup>
  azure_region: <base64_encoded_azure_region>

```



IMPORTANT

Before upgrading a cluster that uses manually maintained credentials, you must ensure that the CCO is in an upgradeable state.

7.8.9.2. Configuring an Azure cluster to use short-term credentials

To install a cluster that uses Microsoft Entra Workload ID, you must configure the Cloud Credential Operator utility and create the required Azure resources for your cluster.

7.8.9.2.1. Configuring the Cloud Credential Operator utility

To create and manage cloud credentials from outside of the cluster when the Cloud Credential Operator (CCO) is operating in manual mode, extract and prepare the CCO utility (**ccoctl**) binary.



NOTE

The **ccoctl** utility is a Linux binary that must run in a Linux environment.

Prerequisites

- You have access to an OpenShift Container Platform account with cluster administrator access.
- You have installed the OpenShift CLI (**oc**).

- You have created a global Microsoft Azure account for the **ccoctl** utility to use with the following permissions:

Example 7.35. Required Azure permissions

- Microsoft.Resources/subscriptions/resourceGroups/read
- Microsoft.Resources/subscriptions/resourceGroups/write
- Microsoft.Resources/subscriptions/resourceGroups/delete
- Microsoft.Authorization/roleAssignments/read
- Microsoft.Authorization/roleAssignments/delete
- Microsoft.Authorization/roleAssignments/write
- Microsoft.Authorization/roleDefinitions/read
- Microsoft.Authorization/roleDefinitions/write
- Microsoft.Authorization/roleDefinitions/delete
- Microsoft.Storage/storageAccounts/listkeys/action
- Microsoft.Storage/storageAccounts/delete
- Microsoft.Storage/storageAccounts/read
- Microsoft.Storage/storageAccounts/write
- Microsoft.Storage/storageAccounts/blobServices/containers/write
- Microsoft.Storage/storageAccounts/blobServices/containers/delete
- Microsoft.Storage/storageAccounts/blobServices/containers/read
- Microsoft.ManagedIdentity/userAssignedIdentities/delete
- Microsoft.ManagedIdentity/userAssignedIdentities/read
- Microsoft.ManagedIdentity/userAssignedIdentities/write
- Microsoft.ManagedIdentity/userAssignedIdentities/federatedIdentityCredentials/read
- Microsoft.ManagedIdentity/userAssignedIdentities/federatedIdentityCredentials/write
- Microsoft.ManagedIdentity/userAssignedIdentities/federatedIdentityCredentials/delete
- Microsoft.Storage/register/action
- Microsoft.ManagedIdentity/register/action

Procedure

1. Set a variable for the OpenShift Container Platform release image by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

- Obtain the CCO container image from the OpenShift Container Platform release image by running the following command:

```
$ CCO_IMAGE=$(oc adm release info --image-for='cloud-credential-operator'
$RELEASE_IMAGE -a ~/.pull-secret)
```



NOTE

Ensure that the architecture of the **\$RELEASE_IMAGE** matches the architecture of the environment in which you will use the **ccoctl** tool.

- Extract the **ccoctl** binary from the CCO container image within the OpenShift Container Platform release image by running the following command:

```
$ oc image extract $CCO_IMAGE \
--file="/usr/bin/ccoctl.<rhel_version>" \
-a ~/.pull-secret
```

- For **<rhel_version>**, specify the value that corresponds to the version of Red Hat Enterprise Linux (RHEL) that the host uses. If no value is specified, **ccoctl.rhel8** is used by default. The following values are valid:

- rhel8**: Specify this value for hosts that use RHEL 8.
- rhel9**: Specify this value for hosts that use RHEL 9.

- Change the permissions to make **ccoctl** executable by running the following command:

```
$ chmod 775 ccoctl.<rhel_version>
```

Verification

- To verify that **ccoctl** is ready to use, display the help file by running the following command:

```
$ ccoctl --help
```

Output of **ccoctl --help**

```
OpenShift credentials provisioning tool
```

```
Usage:
ccoctl [command]
```

```
Available Commands:
```

```
aws      Manage credentials objects for AWS cloud
azure    Manage credentials objects for Azure
gcp      Manage credentials objects for Google cloud
help     Help about any command
ibmcloud Manage credentials objects for IBM Cloud
nutanix  Manage credentials objects for Nutanix
```

Flags:

-h, --help help for ccoctl

Use "ccoctl [command] --help" for more information about a command.

7.8.9.2.2. Creating Azure resources with the Cloud Credential Operator utility

You can use the **ccoctl azure create-all** command to automate the creation of Azure resources.

**NOTE**

By default, **ccoctl** creates objects in the directory in which the commands are run. To create the objects in a different directory, use the **--output-dir** flag. This procedure uses **<path_to_ccoctl_output_dir>** to refer to this directory.

Prerequisites

You must have:

- Extracted and prepared the **ccoctl** binary.
- Access to your Microsoft Azure account by using the Azure CLI.

Procedure

1. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

2. Extract the list of **CredentialsRequest** objects from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included 1 \
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml 2 \
  --to=<path_to_directory_for_credentials_requests> 3
```

- 1** The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2** Specify the location of the **install-config.yaml** file.
- 3** Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

**NOTE**

This command might take a few moments to run.


- To enable the **ccoctl** utility to detect your Azure credentials automatically, log in to the Azure CLI by running the following command:

```
$ az login
```

- Use the **ccoctl** tool to process all **CredentialsRequest** objects by running the following command:

```
$ ccoctl azure create-all \
  --name=<azure_infra_name> \ 1
  --output-dir=<ccoctl_output_dir> \ 2
  --region=<azure_region> \ 3
  --subscription-id=<azure_subscription_id> \ 4
  --credentials-requests-dir=<path_to_credentials_requests_directory> \ 5
  --dnszone-resource-group-name=<azure_dns_zone_resource_group_name> \ 6
  --tenant-id=<azure_tenant_id> \ 7
```

- Specify the user-defined name for all created Azure resources used for tracking.
- Optional: Specify the directory in which you want the **ccoctl** utility to create objects. By default, the utility creates objects in the directory in which the commands are run.
- Specify the Azure region in which cloud resources will be created.
- Specify the Azure subscription ID to use.
- Specify the directory containing the files for the component **CredentialsRequest** objects.
- Specify the name of the resource group containing the cluster's base domain Azure DNS zone.
- Specify the Azure tenant ID to use.



NOTE

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

To see additional optional parameters and explanations of how to use them, run the **azure create-all --help** command.

Verification

- To verify that the OpenShift Container Platform secrets are created, list the files in the **<path_to_ccoctl_output_dir>/manifests** directory:

```
$ ls <path_to_ccoctl_output_dir>/manifests
```

Example output

```
azure-ad-pod-identity-webhook-config.yaml
cluster-authentication-02-config.yaml
```

```

openshift-cloud-controller-manager-azure-cloud-credentials-credentials.yaml
openshift-cloud-network-config-controller-cloud-credentials-credentials.yaml
openshift-cluster-api-capz-manager-bootstrap-credentials-credentials.yaml
openshift-cluster-csi-drivers-azure-disk-credentials-credentials.yaml
openshift-cluster-csi-drivers-azure-file-credentials-credentials.yaml
openshift-image-registry-installer-cloud-credentials-credentials.yaml
openshift-ingress-operator-cloud-credentials-credentials.yaml
openshift-machine-api-azure-cloud-credentials-credentials.yaml

```

You can verify that the Microsoft Entra ID service accounts are created by querying Azure. For more information, refer to Azure documentation on listing Entra ID service accounts.

7.8.9.2.3. Incorporating the Cloud Credential Operator utility manifests

To implement short-term security credentials managed outside the cluster for individual components, you must move the manifest files that the Cloud Credential Operator utility (**ccoctl**) created to the correct directories for the installation program.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have configured the Cloud Credential Operator utility (**ccoctl**).
- You have created the cloud provider resources that are required for your cluster with the **ccoctl** utility.

Procedure

1. If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```

apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...

```

2. If you used the **ccoctl** utility to create a new Azure resource group instead of using an existing resource group, modify the **resourceGroupName** parameter in the **install-config.yaml** as shown:

Sample configuration file snippet

```

apiVersion: v1
baseDomain: example.com
# ...
platform:
  azure:
    resourceGroupName: <azure_infra_name> 1
# ...

```


- 1 This value must match the user-defined name for Azure resources that was specified with the **--name** argument of the **ccoctl azure create-all** command.
3. If you have not previously created installation manifest files, do so by running the following command:


```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.
4. Copy the manifests that the **ccoctl** utility generated to the **manifests** directory that the installation program created by running the following command:


```
$ cp /<path_to_ccoctl_output_dir>/manifests/* ./manifests/
```
5. Copy the private key that the **ccoctl** utility generated in the **tls** directory to the installation directory by running the following command:


```
$ cp -a /<path_to_ccoctl_output_dir>/tls .
```

7.8.10. Optional: Preparing a private Microsoft Azure cluster for a private image registry

By installing a private image registry on a private Microsoft Azure cluster, you can create private storage endpoints. Private storage endpoints disable public facing endpoints to the registry's storage account, adding an extra layer of security to your OpenShift Container Platform deployment. Use the following guide to prepare your private Microsoft Azure cluster for installation with a private image registry.

Prerequisites

- You have access to an OpenShift Container Platform account with cluster administrator access.
- You have installed the OpenShift CLI (oc).
- You have prepared an **install-config.yaml** that includes the following information:
 - The **publish** field is set to **Internal**
- You have set the permissions for creating a private storage endpoint. For more information, see "Azure permissions for installer-provisioned infrastructure".

Procedure

1. If you have not previously created installation manifest files, do so by running the following command:

```
$ ./openshift-install create manifests --dir <installation_directory>
```

This command displays the following messages:

Example output


```
INFO Consuming Install Config from target directory
INFO Manifests created in: <installation_directory>/manifests and
<installation_directory>/openshift
```

2. Create an image registry configuration object and pass in the **networkResourceGroupName**, **subnetName**, and **vnetName** provided by Microsoft Azure. For example:

```
$ touch imageregistry-config.yaml
```

```
apiVersion: imageregistry.operator.openshift.io/v1
kind: Config
metadata:
  name: cluster
spec:
  managementState: "Managed"
  replicas: 2
  rolloutStrategy: RollingUpdate
  storage:
    azure:
      networkAccess:
        internal:
          networkResourceGroupName: <vnet_resource_group> 1
          subnetName: <subnet_name> 2
          vnetName: <vnet_name> 3
      type: Internal
```

- 1** Optional. If you have an existing VNet and subnet setup, replace **<vnet_resource_group>** with the resource group name that contains the existing virtual network (VNet).
- 2** Optional. If you have an existing VNet and subnet setup, replace **<subnet_name>** with the name of the existing compute subnet within the specified resource group.
- 3** Optional. If you have an existing VNet and subnet setup, replace **<vnet_name>** with the name of the existing virtual network (VNet) in the specified resource group.



NOTE

The **imageregistry-config.yaml** file is consumed during the installation process. If desired, you must back it up before installation.

3. Move the **imageregistry-config.yaml** file to the **<installation_directory/manifests>** folder by running the following command:

```
$ mv imageregistry-config.yaml <installation_directory/manifests/>
```

Next steps


- After you have moved the **imageregistry-config.yaml** file to the **<installation_directory/manifests>** folder and set the required permissions, proceed to "Deploying the cluster".

Additional resources

- For the list of permissions needed to create a private storage endpoint, see [Required Azure permissions for installer-provisioned infrastructure](#).

7.8.11. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have an Azure subscription ID and tenant ID.
- If you are installing the cluster using a service principal, you have its application ID and password.
- If you are installing the cluster using a system-assigned managed identity, you have enabled it on the virtual machine that you will run the installation program from.
- If you are installing the cluster using a user-assigned managed identity, you have met these prerequisites:
 - You have its client ID.
 - You have assigned it to the virtual machine that you will run the installation program from.

Procedure

1. Optional: If you have run the installation program on this computer before, and want to use an alternative service principal or managed identity, go to the `~/.azure/` directory and delete the **osServicePrincipal.json** configuration file.
Deleting this file prevents the installation program from automatically reusing subscription and authentication values from a previous installation.
2. Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

- 1 For **<installation_directory>**, specify the location of your customized **./install-config.yaml** file.
- 2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

If the installation program cannot locate the **osServicePrincipal.json** configuration file from a previous installation, you are prompted for Azure subscription and authentication values.

3. Enter the following Azure parameter values for your subscription:
 - **azure subscription id** Enter the subscription ID to use for the cluster.
 - **azure tenant id** Enter the tenant ID.
4. Depending on the Azure identity you are using to deploy the cluster, do one of the following when prompted for the **azure service principal client id**
 - If you are using a service principal, enter its application ID.
 - If you are using a system-assigned managed identity, leave this value blank.
 - If you are using a user-assigned managed identity, specify its client ID.
5. Depending on the Azure identity you are using to deploy the cluster, do one of the following when prompted for the **azure service principal client secret**
 - If you are using a service principal, enter its password.
 - If you are using a system-assigned managed identity, leave this value blank.
 - If you are using a user-assigned managed identity, leave this value blank.

If previously not detected, the installation program creates an **osServicePrincipal.json** configuration file and stores this file in the `~/.azure/` directory on your computer. This ensures that the installation program can load the profile when it is creating an OpenShift Container Platform cluster on the target platform.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to `<installation_directory>/openshift_install.log`.



IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

7.8.12. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

Additional resources

- See [Accessing the web console](#) for more details about accessing and understanding the OpenShift Container Platform web console.

7.8.13. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service

7.8.14. Next steps

- [Customize your cluster](#).
- If necessary, you can [opt out of remote health reporting](#).

7.9. INSTALLING A CLUSTER ON AZURE INTO A GOVERNMENT REGION

In OpenShift Container Platform version 4.16, you can install a cluster on Microsoft Azure into a government region. To configure the government region, you modify parameters in the **install-config.yaml** file before you install the cluster.

7.9.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [configured an Azure account](#) to host the cluster and determined the tested and validated government region to deploy the cluster to.
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.
- If the cloud identity and access management (IAM) APIs are not accessible in your environment, or if you do not want to store an administrator-level credential secret in the **kube-system** namespace, you can [manually create and maintain long-term credentials](#).
- If you use customer-managed encryption keys, you [prepared your Azure environment for encryption](#).

7.9.2. Azure government regions

OpenShift Container Platform supports deploying a cluster to [Microsoft Azure Government \(MAG\)](#) regions. MAG is specifically designed for US government agencies at the federal, state, and local level, as well as contractors, educational institutions, and other US customers that must run sensitive workloads on Azure. MAG is composed of government-only data center regions, all granted an [Impact Level 5 Provisional Authorization](#).

Installing to a MAG region requires manually configuring the Azure Government dedicated cloud instance and region in the **install-config.yaml** file. You must also update your service principal to reference the appropriate government environment.




NOTE

The Azure government region cannot be selected using the guided terminal prompts from the installation program. You must define the region manually in the **install-config.yaml** file. Remember to also set the dedicated cloud instance, like **AzureUSGovernmentCloud**, based on the region specified.

7.9.3. Private clusters

You can deploy a private OpenShift Container Platform cluster that does not expose external endpoints. Private clusters are accessible from only an internal network and are not visible to the internet.

By default, OpenShift Container Platform is provisioned to use publicly-accessible DNS and endpoints. A private cluster sets the DNS, Ingress Controller, and API server to private when you deploy your cluster. This means that the cluster resources are only accessible from your internal network and are not visible to the internet.



IMPORTANT

If the cluster has any public subnets, load balancer services created by administrators might be publicly accessible. To ensure cluster security, verify that these services are explicitly annotated as private.

To deploy a private cluster, you must:

- Use existing networking that meets your requirements. Your cluster resources might be shared between other clusters on the network.
- Deploy from a machine that has access to:
 - The API services for the cloud to which you provision.
 - The hosts on the network that you provision.
 - The internet to obtain installation media.

You can use any machine that meets these access requirements and follows your company's guidelines. For example, this machine can be a bastion host on your cloud network or a machine that has access to the network through a VPN.

7.9.3.1. Private clusters in Azure

To create a private cluster on Microsoft Azure, you must provide an existing private VNet and subnets to host the cluster. The installation program must also be able to resolve the DNS records that the cluster requires. The installation program configures the Ingress Operator and API server for only internal traffic.

Depending how your network connects to the private VNET, you might need to use a DNS forwarder to resolve the cluster's private DNS records. The cluster's machines use **168.63.129.16** internally for DNS resolution. For more information, see [What is Azure Private DNS?](#) and [What is IP address 168.63.129.16?](#)

in the Azure documentation.

The cluster still requires access to internet to access the Azure APIs.

The following items are not required or created when you install a private cluster:

- A **BaseDomainResourceGroup**, since the cluster does not create public records
- Public IP addresses
- Public DNS records
- Public endpoints

The cluster is configured so that the Operators do not create public records for the cluster and all cluster machines are placed in the private subnets that you specify.

7.9.3.1.1. Limitations

Private clusters on Azure are subject to only the limitations that are associated with the use of an existing VNet.

7.9.3.2. User-defined outbound routing

In OpenShift Container Platform, you can choose your own outbound routing for a cluster to connect to the internet. This allows you to skip the creation of public IP addresses and the public load balancer.

You can configure user-defined routing by modifying parameters in the **install-config.yaml** file before installing your cluster. A pre-existing VNet is required to use outbound routing when installing a cluster; the installation program is not responsible for configuring this.

When configuring a cluster to use user-defined routing, the installation program does not create the following resources:

- Outbound rules for access to the internet.
- Public IPs for the public load balancer.
- Kubernetes Service object to add the cluster machines to the public load balancer for outbound requests.

You must ensure the following items are available before setting user-defined routing:

- Egress to the internet is possible to pull container images, unless using an OpenShift image registry mirror.
- The cluster can access Azure APIs.
- Various allowlist endpoints are configured. You can reference these endpoints in the *Configuring your firewall* section.

There are several pre-existing networking setups that are supported for internet access using user-defined routing.

7.9.4. About reusing a VNet for your OpenShift Container Platform cluster

In OpenShift Container Platform 4.16, you can deploy a cluster into an existing Azure Virtual Network (VNet) in Microsoft Azure. If you do, you must also use existing subnets within the VNet and routing rules.

By deploying OpenShift Container Platform into an existing Azure VNet, you might be able to avoid service limit constraints in new accounts or more easily abide by the operational constraints that your company's guidelines set. This is a good option to use if you cannot obtain the infrastructure creation permissions that are required to create the VNet.

7.9.4.1. Requirements for using your VNet

When you deploy a cluster by using an existing VNet, you must perform additional network configuration before you install the cluster. In installer-provisioned infrastructure clusters, the installer usually creates the following components, but it does not create them when you install into an existing VNet:

- Subnets
- Route tables
- VNets
- Network Security Groups



NOTE

The installation program requires that you use the cloud-provided DNS server. Using a custom DNS server is not supported and causes the installation to fail.

If you use a custom VNet, you must correctly configure it and its subnets for the installation program and the cluster to use. The installation program cannot subdivide network ranges for the cluster to use, set route tables for the subnets, or set VNet options like DHCP, so you must do so before you install the cluster.

The cluster must be able to access the resource group that contains the existing VNet and subnets. While all of the resources that the cluster creates are placed in a separate resource group that it creates, some network resources are used from a separate group. Some cluster Operators must be able to access resources in both resource groups. For example, the Machine API controller attaches NICS for the virtual machines that it creates to subnets from the networking resource group.

Your VNet must meet the following characteristics:

- The VNet's CIDR block must contain the **Networking.MachineCIDR** range, which is the IP address pool for cluster machines.
- The VNet and its subnets must belong to the same resource group, and the subnets must be configured to use Azure-assigned DHCP IP addresses instead of static IP addresses.

You must provide two subnets within your VNet, one for the control plane machines and one for the compute machines. Because Azure distributes machines in different availability zones within the region that you specify, your cluster will have high availability by default.

**NOTE**

By default, if you specify availability zones in the **install-config.yaml** file, the installation program distributes the control plane machines and the compute machines across [these availability zones](#) within [a region](#). To ensure high availability for your cluster, select a region with at least three availability zones. If your region contains fewer than three availability zones, the installation program places more than one control plane machine in the available zones.

To ensure that the subnets that you provide are suitable, the installation program confirms the following data:


- All the specified subnets exist.
- There are two private subnets, one for the control plane machines and one for the compute machines.
- The subnet CIDRs belong to the machine CIDR that you specified. Machines are not provisioned in availability zones that you do not provide private subnets for. If required, the installation program creates public load balancers that manage the control plane and worker nodes, and Azure allocates a public IP address to them.

**NOTE**

If you destroy a cluster that uses an existing VNet, the VNet is not deleted.

7.9.4.1.1. Network security group requirements

The network security groups for the subnets that host the compute and control plane machines require specific access to ensure that the cluster communication is correct. You must create rules to allow access to the required cluster communication ports.

**IMPORTANT**

The network security group rules must be in place before you install the cluster. If you attempt to install a cluster without the required access, the installation program cannot reach the Azure APIs, and installation fails.

Table 7.18. Required ports

Port	Description	Control plane	Compute
80	Allows HTTP traffic		x
443	Allows HTTPS traffic		x
6443	Allows communication to the control plane machines	x	
22623	Allows internal communication to the machine config server for provisioning machines	x	

1. If you are using Azure Firewall to restrict the internet access, then [you can configure Azure Firewall to allow the Azure APIs](#). A network security group rule is not needed.



IMPORTANT

Currently, there is no supported way to block or restrict the machine config server endpoint. The machine config server must be exposed to the network so that newly-provisioned machines, which have no existing configuration or state, are able to fetch their configuration. In this model, the root of trust is the certificate signing requests (CSR) endpoint, which is where the kubelet sends its certificate signing request for approval to join the cluster. Because of this, machine configs should not be used to distribute sensitive information, such as secrets and certificates.

To ensure that the machine config server endpoints, ports 22623 and 22624, are secured in bare metal scenarios, customers must configure proper network policies.

Because cluster components do not modify the user-provided network security groups, which the Kubernetes controllers update, a pseudo-network security group is created for the Kubernetes controller to modify without impacting the rest of the environment.

Additional resources

- [About the OpenShift SDN network plugin](#)
- [Configuring your firewall](#)

7.9.4.2. Division of permissions

Starting with OpenShift Container Platform 4.3, you do not need all of the permissions that are required for an installation program-provisioned infrastructure cluster to deploy a cluster. This change mimics the division of permissions that you might have at your company: some individuals can create different resources in your clouds than others. For example, you might be able to create application-specific items, like instances, storage, and load balancers, but not networking-related components such as VNets, subnet, or ingress rules.

The Azure credentials that you use when you create your cluster do not need the networking permissions that are required to make VNets and core networking components within the VNet, such as subnets, routing tables, internet gateways, NAT, and VPN. You still need permission to make the application resources that the machines within the cluster require, such as load balancers, security groups, storage accounts, and nodes.

7.9.4.3. Isolation between clusters

Because the cluster is unable to modify network security groups in an existing subnet, there is no way to isolate clusters from each other on the VNet.

7.9.5. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.

- Access [Quay.io](https://quay.io) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.



IMPORTANT

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

7.9.6. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.



NOTE


You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.

**NOTE**

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.


2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.

**NOTE**

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```

**NOTE**

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

7.9.7. Obtaining the installation program


Before you install OpenShift Container Platform, download the installation file on the host you are using for installation.

Prerequisites

- You have a computer that runs Linux or macOS, with 500 MB of local disk space.

Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select your infrastructure provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.



IMPORTANT

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.



IMPORTANT

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.

4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#) . This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

7.9.8. Manually creating the installation configuration file

Installing the cluster requires that you manually create the installation configuration file.

Prerequisites

- You have an SSH public key on your local machine to provide to the installation program. The key will be used for SSH authentication onto your cluster nodes for debugging and disaster recovery.
- You have obtained the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create an installation directory to store your required installation assets in:

```
$ mkdir <installation_directory>
```



IMPORTANT

You must create a directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.


2. Customize the sample **install-config.yaml** file template that is provided and save it in the **<installation_directory>**.



NOTE

You must name this configuration file **install-config.yaml**.

3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the next step of the installation process. You must back it up now.

Additional resources

- [Installation configuration parameters for Azure](#)

7.9.8.1. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 7.19. Minimum resource requirements

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Bootstrap	RHCOS	4	16 GB	100 GB	300

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Control plane	RHCOS	4	16 GB	100 GB	300
Compute	RHCOS, RHEL 8.6 and later [3]	2	8 GB	100 GB	300

1. One vCPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio: (threads per core × cores) × sockets = vCPUs.
2. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes which require a 10 ms p99 fsync duration. Note that on many cloud platforms, storage size and IOPS scale together, so you might need to over-allocate storage volume to obtain sufficient performance.
3. As with all user-provisioned installations, if you choose to use RHEL compute machines in your cluster, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Use of RHEL 7 compute machines is deprecated and has been removed in OpenShift Container Platform 4.10 and later.



NOTE

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).



IMPORTANT

You are required to use Azure virtual machines that have the **premiumIO** parameter set to **true**.

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

7.9.8.2. Tested instance types for Azure

The following Microsoft Azure instance types have been tested with OpenShift Container Platform.

Example 7.36. Machine types based on 64-bit x86 architecture

- **c4.***
- **c5.***
- **c5a.***
- **i3.***
- **m4.***
- **m5.***
- **m5a.***
- **m6a.***
- **m6i.***
- **r4.***
- **r5.***
- **r5a.***
- **r6i.***
- **t3.***
- **t3a.***

7.9.8.3. Enabling trusted launch for Azure VMs

You can enable two trusted launch features when installing your cluster on Azure: [secure boot](#) and [virtualized Trusted Platform Modules](#).

See the Azure documentation about [virtual machine sizes](#) to learn what sizes of virtual machines support these features.



IMPORTANT

Trusted launch is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

Prerequisites

- You have created an **install-config.yaml** file.

Procedure

- Use a text editor to edit the **install-config.yaml** file prior to deploying your cluster and add the following stanza:

```
controlPlane: ❶
  platform:
    azure:
      settings:
        securityType: TrustedLaunch ❷
        trustedLaunch:
          uefiSettings:
            secureBoot: Enabled ❸
            virtualizedTrustedPlatformModule: Enabled ❹
```

- Specify **controlPlane.platform.azure** or **compute.platform.azure** to enable trusted launch on only control plane or compute nodes respectively. Specify **platform.azure.defaultMachinePlatform** to enable trusted launch on all nodes.
- Enable trusted launch features.
- Enable secure boot. For more information, see the Azure documentation about [secure boot](#).
- Enable the virtualized Trusted Platform Module. For more information, see the Azure documentation about [virtualized Trusted Platform Modules](#).

7.9.8.4. Enabling confidential VMs

You can enable confidential VMs when installing your cluster. You can enable confidential VMs for compute nodes, control plane nodes, or all nodes.



IMPORTANT

Using confidential VMs is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

You can use confidential VMs with the following VM sizes:

- DCasv5-series
- DCadsv5-series
- ECasv5-series

- ECadsV5-series



IMPORTANT

Confidential VMs are currently not supported on 64-bit ARM architectures.

Prerequisites

- You have created an **install-config.yaml** file.

Procedure

- Use a text editor to edit the **install-config.yaml** file prior to deploying your cluster and add the following stanza:

```
controlPlane: 1
platform:
  azure:
    settings:
      securityType: ConfidentialVM 2
      confidentialVM:
        uefiSettings:
          secureBoot: Enabled 3
          virtualizedTrustedPlatformModule: Enabled 4
    osDisk:
      securityProfile:
        securityEncryptionType: VMGuestStateOnly 5
```

- 1 Specify **controlPlane.platform.azure** or **compute.platform.azure** to deploy confidential VMs on only control plane or compute nodes respectively. Specify **platform.azure.defaultMachinePlatform** to deploy confidential VMs on all nodes.
- 2 Enable confidential VMs.
- 3 Enable secure boot. For more information, see the Azure documentation about [secure boot](#).
- 4 Enable the virtualized Trusted Platform Module. For more information, see the Azure documentation about [virtualized Trusted Platform Modules](#).
- 5 Specify **VMGuestStateOnly** to encrypt the VM guest state.

7.9.8.5. Sample customized install-config.yaml file for Azure

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.



IMPORTANT

This sample YAML file is provided for reference only. You must obtain your **install-config.yaml** file by using the installation program and modify it.

```
apiVersion: v1
baseDomain: example.com 1
controlPlane: 2
  hyperthreading: Enabled 3 4
  name: master
  platform:
    azure:
      encryptionAtHost: true
      ultraSSDCapability: Enabled
    osDisk:
      diskSizeGB: 1024 5
      diskType: Premium_LRS
      diskEncryptionSet:
        resourceGroup: disk_encryption_set_resource_group
        name: disk_encryption_set_name
        subscriptionId: secondary_subscription_id
    osImage:
      publisher: example_publisher_name
      offer: example_image_offer
      sku: example_offer_sku
      version: example_image_version
      type: Standard_D8s_v3
  replicas: 3
compute: 6
- hyperthreading: Enabled 7
  name: worker
  platform:
    azure:
      ultraSSDCapability: Enabled
      type: Standard_D2s_v3
      encryptionAtHost: true
    osDisk:
      diskSizeGB: 512 8
      diskType: Standard_LRS
      diskEncryptionSet:
        resourceGroup: disk_encryption_set_resource_group
        name: disk_encryption_set_name
        subscriptionId: secondary_subscription_id
    osImage:
      publisher: example_publisher_name
      offer: example_image_offer
      sku: example_offer_sku
      version: example_image_version
  zones: 9
  - "1"
  - "2"
  - "3"
  replicas: 5
metadata:
  name: test-cluster 10
networking:
  clusterNetwork:
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  machineNetwork:
```

```

- cidr: 10.0.0.0/16
networkType: OVNKubernetes 11
serviceNetwork:
- 172.30.0.0/16
platform:
  azure:
    defaultMachinePlatform:
      osImage: 12
      publisher: example_publisher_name
      offer: example_image_offer
      sku: example_offer_sku
      version: example_image_version
      ultraSSDCapability: Enabled
    baseDomainResourceGroupName: resource_group 13
    region: usgovvirginia
    resourceGroupName: existing_resource_group 14
    networkResourceGroupName: vnet_resource_group 15
    virtualNetwork: vnet 16
    controlPlaneSubnet: control_plane_subnet 17
    computeSubnet: compute_subnet 18
    outboundType: UserDefinedRouting 19
    cloudName: AzureUSGovernmentCloud 20
  pullSecret: '{"auths": ...}' 21
  fips: false 22
  sshKey: ssh-ed25519 AAAA... 23
  publish: Internal 24

```

1 **10** **21** Required.

2 **6** If you do not provide these parameters and values, the installation program provides the default value.

3 **7** The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.

4 Whether to enable or disable simultaneous multithreading, or **hyperthreading**. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores. You can disable it by setting the parameter value to **Disabled**. If you disable simultaneous multithreading in some cluster machines, you must disable it in all cluster machines.



IMPORTANT

If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance. Use larger virtual machine types, such as **Standard_D8s_v3**, for your machines if you disable simultaneous multithreading.

5 **8** You can specify the size of the disk to use in GB. Minimum recommendation for control plane nodes is 1024 GB.

9 Specify a list of zones to deploy your machines to. For high availability, specify at least two zones.

- 11 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 12 Optional: A custom Red Hat Enterprise Linux CoreOS (RHCOS) image that should be used to boot control plane and compute machines. The **publisher**, **offer**, **sku**, and **version** parameters under **platform.azure.defaultMachinePlatform.osImage** apply to both control plane and compute machines. If the parameters under **controlPlane.platform.azure.osImage** or **compute.platform.azure.osImage** are set, they override the **platform.azure.defaultMachinePlatform.osImage** parameters.
- 13 Specify the name of the resource group that contains the DNS zone for your base domain.
- 14 Specify the name of an already existing resource group to install your cluster to. If undefined, a new resource group is created for the cluster.
- 15 If you use an existing VNet, specify the name of the resource group that contains it.
- 16 If you use an existing VNet, specify its name.
- 17 If you use an existing VNet, specify the name of the subnet to host the control plane machines.
- 18 If you use an existing VNet, specify the name of the subnet to host the compute machines.
- 19 You can customize your own outbound routing. Configuring user-defined routing prevents exposing external endpoints in your cluster. User-defined routing for egress requires deploying your cluster to an existing VNet.
- 20 Specify the name of the Azure cloud environment to deploy your cluster to. Set **AzureUSGovernmentCloud** to deploy to a Microsoft Azure Government (MAG) region. The default value is **AzurePublicCloud**.
- 22 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.



IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 23 You can optionally provide the **sshKey** value that you use to access the machines in your cluster.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- 24 How to publish the user-facing endpoints of your cluster. Set **publish** to **Internal** to deploy a private cluster, which cannot be accessed from the internet. The default value is **External**.

7.9.8.6. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure


- Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- A proxy URL to use for creating HTTPS connections outside the cluster.
- A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all

destinations.

- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.



NOTE

The installation program does not support the proxy **readinessEndpoints** field.



NOTE

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.



NOTE


Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

Additional resources

- For more details about Accelerated Networking, see [Accelerated Networking for Microsoft Azure VMs](#).

7.9.9. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have an Azure subscription ID and tenant ID.
- If you are installing the cluster using a service principal, you have its application ID and password.
- If you are installing the cluster using a system-assigned managed identity, you have enabled it on the virtual machine that you will run the installation program from.
- If you are installing the cluster using a user-assigned managed identity, you have met these prerequisites:
 - You have its client ID.
 - You have assigned it to the virtual machine that you will run the installation program from.

Procedure

1. Optional: If you have run the installation program on this computer before, and want to use an alternative service principal or managed identity, go to the `~/.azure/` directory and delete the **osServicePrincipal.json** configuration file.
Deleting this file prevents the installation program from automatically reusing subscription and authentication values from a previous installation.
2. Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

- 1** For `<installation_directory>`, specify the location of your customized `./install-config.yaml` file.
- 2** To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

If the installation program cannot locate the **osServicePrincipal.json** configuration file from a previous installation, you are prompted for Azure subscription and authentication values.

3. Enter the following Azure parameter values for your subscription:
 - **azure subscription id** Enter the subscription ID to use for the cluster.
 - **azure tenant id** Enter the tenant ID.
4. Depending on the Azure identity you are using to deploy the cluster, do one of the following when prompted for the **azure service principal client id**
 - If you are using a service principal, enter its application ID.
 - If you are using a system-assigned managed identity, leave this value blank.

- If you are using a user-assigned managed identity, specify its client ID.
5. Depending on the Azure identity you are using to deploy the cluster, do one of the following when prompted for the **azure service principal client secret**
- If you are using a service principal, enter its password.
 - If you are using a system-assigned managed identity, leave this value blank.
 - If you are using a user-assigned managed identity, leave this value blank.

If previously not detected, the installation program creates an **osServicePrincipal.json** configuration file and stores this file in the `~/.azure/` directory on your computer. This ensures that the installation program can load the profile when it is creating an OpenShift Container Platform cluster on the target platform.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to `<installation_directory>/openshift_install.log`.




IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

7.9.10. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.



IMPORTANT

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**. To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

-

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.



NOTE

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your PATH.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

7.9.11. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

Additional resources

- See [Accessing the web console](#) for more details about accessing and understanding the OpenShift Container Platform web console.

7.9.12. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service

7.9.13. Next steps

- [Customize your cluster.](#)
- If necessary, you can [opt out of remote health reporting](#) .

7.10. INSTALLING A CLUSTER ON AZURE IN A RESTRICTED NETWORK WITH USER-PROVISIONED INFRASTRUCTURE

In OpenShift Container Platform, you can install a cluster on Microsoft Azure by using infrastructure that you provide.

Several [Azure Resource Manager](#) (ARM) templates are provided to assist in completing these steps or to help model your own.




IMPORTANT

The steps for performing a user-provisioned infrastructure installation are provided as an example only. Installing a cluster with infrastructure you provide requires knowledge of the cloud provider and the installation process of OpenShift Container Platform. Several ARM templates are provided to assist in completing these steps or to help model your own. You are also free to create the required resources through other methods.

Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [configured an Azure account](#) to host the cluster and determined the tested and validated region to deploy the cluster to.
- You [mirrored the images for a disconnected installation](#) to your registry and obtained the **imageContentSources** data for your version of OpenShift Container Platform.



IMPORTANT

Because the installation media is on the mirror host, you must use that computer to complete all installation steps.

- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.
- If the cloud identity and access management (IAM) APIs are not accessible in your environment, or if you do not want to store an administrator-level credential secret in the **kube-system** namespace, you have [manually created long-term credentials](#).
- If you use customer-managed encryption keys, you [prepared your Azure environment for encryption](#).

7.10.1. About installations in restricted networks

In OpenShift Container Platform 4.16, you can perform an installation that does not require an active

connection to the internet to obtain software components. Restricted network installations can be completed using installer-provisioned infrastructure or user-provisioned infrastructure, depending on the cloud platform to which you are installing the cluster.

If you choose to perform a restricted network installation on a cloud platform, you still require access to its cloud APIs. Some cloud functions, like Amazon Web Service's Route 53 DNS and IAM services, require internet access. Depending on your network, you might require less internet access for an installation on bare metal hardware, Nutanix, or on VMware vSphere.

To complete a restricted network installation, you must create a registry that mirrors the contents of the OpenShift image registry and contains the installation media. You can create this registry on a mirror host, which can access both the internet and your closed network, or by using other methods that meet your restrictions.



IMPORTANT

Because of the complexity of the configuration for user-provisioned installations, consider completing a standard user-provisioned infrastructure installation before you attempt a restricted network installation using user-provisioned infrastructure. Completing this test installation might make it easier to isolate and troubleshoot any issues that might arise during your installation in a restricted network.

7.10.1.1. Additional limits

Clusters in restricted networks have the following additional limitations and restrictions:

- The **ClusterVersion** status includes an **Unable to retrieve available updates** error.
- By default, you cannot use the contents of the Developer Catalog because you cannot access the required image stream tags.

7.10.1.2. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.




IMPORTANT

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

7.10.2. Configuring your Azure project

Before you can install OpenShift Container Platform, you must configure an Azure project to host it.




IMPORTANT

All Azure resources that are available through public endpoints are subject to resource name restrictions, and you cannot create resources that use certain terms. For a list of terms that Azure restricts, see [Resolve reserved resource name errors](#) in the Azure documentation.

7.10.2.1. Azure account limits

The OpenShift Container Platform cluster uses a number of Microsoft Azure components, and the default [Azure subscription and service limits, quotas, and constraints](#) affect your ability to install OpenShift Container Platform clusters.



IMPORTANT

Default limits vary by offer category types, such as Free Trial and Pay-As-You-Go, and by series, such as Dv2, F, and G. For example, the default for Enterprise Agreement subscriptions is 350 cores.


Check the limits for your subscription type and if necessary, increase quota limits for your account before you install a default cluster on Azure.

The following table summarizes the Azure components whose limits can impact your ability to install and run OpenShift Container Platform clusters.

Component	Number of components required by default	Default Azure limit	Description

Component	Number of components required by default	Default Azure limit	Description
vCPU	44	20 per region	<p>A default cluster requires 44 vCPUs, so you must increase the account limit.</p> <p>By default, each cluster creates the following instances:</p> <ul style="list-style-type: none"> • One bootstrap machine, which is removed after installation • Three control plane machines • Three compute machines <p>Because the bootstrap and control plane machines use Standard_D8s_v3 virtual machines, which use 8 vCPUs, and the compute machines use Standard_D4s_v3 virtual machines, which use 4 vCPUs, a default cluster requires 44 vCPUs. The bootstrap node VM, which uses 8 vCPUs, is used only during installation.</p> <p>To deploy more worker nodes, enable autoscaling, deploy large workloads, or use a different instance type, you must further increase the vCPU limit for your account to ensure that your cluster can deploy the machines that you require.</p>
OS Disk	7		<p>Each cluster machine must have a minimum of 100 GB of storage and 300 IOPS. While these are the minimum supported values, faster storage is recommended for production clusters and clusters with intensive workloads. For more information about optimizing storage for performance, see the page titled "Optimizing storage" in the "Scalability and performance" section.</p>
VNet	1	1000 per region	<p>Each default cluster requires one Virtual Network (VNet), which contains two subnets.</p>
Network interfaces	7	65,536 per region	<p>Each default cluster requires seven network interfaces. If you create more machines or your deployed workloads create load balancers, your cluster uses more network interfaces.</p>

Component	Number of components required by default	Default Azure limit	Description						
Network security groups	2	5000	<p>Each cluster creates network security groups for each subnet in the VNet. The default cluster creates network security groups for the control plane and for the compute node subnets:</p> <table border="1"> <tr> <td>control plane</td> <td>Allows the control plane machines to be reached on port 6443 from anywhere</td> </tr> <tr> <td>node</td> <td>Allows worker nodes to be reached from the internet on ports 80 and 443</td> </tr> </table>	control plane	Allows the control plane machines to be reached on port 6443 from anywhere	node	Allows worker nodes to be reached from the internet on ports 80 and 443		
control plane	Allows the control plane machines to be reached on port 6443 from anywhere								
node	Allows worker nodes to be reached from the internet on ports 80 and 443								
Network load balancers	3	1000 per region	<p>Each cluster creates the following load balancers:</p> <table border="1"> <tr> <td>default</td> <td>Public IP address that load balances requests to ports 80 and 443 across worker machines</td> </tr> <tr> <td>internal</td> <td>Private IP address that load balances requests to ports 6443 and 22623 across control plane machines</td> </tr> <tr> <td>external</td> <td>Public IP address that load balances requests to port 6443 across control plane machines</td> </tr> </table> <p>If your applications create more Kubernetes LoadBalancer service objects, your cluster uses more load balancers.</p>	default	Public IP address that load balances requests to ports 80 and 443 across worker machines	internal	Private IP address that load balances requests to ports 6443 and 22623 across control plane machines	external	Public IP address that load balances requests to port 6443 across control plane machines
default	Public IP address that load balances requests to ports 80 and 443 across worker machines								
internal	Private IP address that load balances requests to ports 6443 and 22623 across control plane machines								
external	Public IP address that load balances requests to port 6443 across control plane machines								
Public IP addresses	3		Each of the two public load balancers uses a public IP address. The bootstrap machine also uses a public IP address so that you can SSH into the machine to troubleshoot issues during installation. The IP address for the bootstrap node is used only during installation.						
Private IP addresses	7		The internal load balancer, each of the three control plane machines, and each of the three worker machines each use a private IP address.						

Component	Number of components required by default	Default Azure limit	Description
Spot VM vCPUs (optional)	0 If you configure spot VMs, your cluster must have two spot VM vCPUs for every compute node.	20 per region	<p>This is an optional component. To use spot VMs, you must increase the Azure default limit to at least twice the number of compute nodes in your cluster.</p> <div style="display: flex; align-items: center;">  <div> <p>NOTE</p> <p>Using spot VMs for control plane nodes is not recommended.</p> </div> </div>

Additional resources


- [Optimizing storage](#)

7.10.2.2. Configuring a public DNS zone in Azure

To install OpenShift Container Platform, the Microsoft Azure account you use must have a dedicated public hosted DNS zone in your account. This zone must be authoritative for the domain. This service provides cluster DNS resolution and name lookup for external connections to the cluster.

Procedure

1. Identify your domain, or subdomain, and registrar. You can transfer an existing domain and registrar or obtain a new one through Azure or another source.



NOTE

For more information about purchasing domains through Azure, see [Buy a custom domain name for Azure App Service](#) in the Azure documentation.

2. If you are using an existing domain and registrar, migrate its DNS to Azure. See [Migrate an active DNS name to Azure App Service](#) in the Azure documentation.
3. Configure DNS for your domain. Follow the steps in the [Tutorial: Host your domain in Azure DNS](#) in the Azure documentation to create a public hosted zone for your domain or subdomain, extract the new authoritative name servers, and update the registrar records for the name servers that your domain uses.
Use an appropriate root domain, such as **openshiftcorp.com**, or subdomain, such as **clusters.openshiftcorp.com**.
4. If you use a subdomain, follow your company's procedures to add its delegation records to the parent domain.

You can view Azure's DNS solution by visiting this [example for creating DNS zones](#).

7.10.2.3. Increasing Azure account limits

To increase an account limit, file a support request on the Azure portal.

**NOTE**

You can increase only one type of quota per support request.

Procedure

1. From the Azure portal, click **Help + support** in the lower left corner.
2. Click **New support request** and then select the required values:
 - a. From the **Issue type** list, select **Service and subscription limits (quotas)**
 - b. From the **Subscription** list, select the subscription to modify.
 - c. From the **Quota type** list, select the quota to increase. For example, select **Compute-VM (cores-vCPUs) subscription limit increases** to increase the number of vCPUs, which is required to install a cluster.
 - d. Click **Next: Solutions**.
3. On the **Problem Details** page, provide the required information for your quota increase:
 - a. Click **Provide details** and provide the required details in the **Quota details** window.
 - b. In the SUPPORT METHOD and CONTACT INFO sections, provide the issue severity and your contact details.
4. Click **Next: Review + create** and then click **Create**.

7.10.2.4. Certificate signing requests management

Because your cluster has limited access to automatic machine management when you use infrastructure that you provision, you must provide a mechanism for approving cluster certificate signing requests (CSRs) after installation. The **kube-controller-manager** only approves the kubelet client CSRs. The **machine-approver** cannot guarantee the validity of a serving certificate that is requested by using kubelet credentials because it cannot confirm that the correct machine issued the request. You must determine and implement a method of verifying the validity of the kubelet serving certificate requests and approving them.

7.10.2.5. Required Azure roles

An OpenShift Container Platform cluster requires an Azure identity to create and manage Azure resources. Before you create the identity, verify that your environment meets the following requirements:

- The Azure account that you use to create the identity is assigned the **User Access Administrator** and **Contributor** roles. These roles are required when:
 - Creating a service principal or user-assigned managed identity.
 - Enabling a system-assigned managed identity on a virtual machine.
- If you are going to use a service principal to complete the installation, verify that the Azure account that you use to create the identity is assigned the **microsoft.directory/servicePrincipals/createAsOwner** permission in Microsoft Entra ID.

To set roles on the Azure portal, see the [Manage access to Azure resources using RBAC and the Azure portal](#) in the Azure documentation.

7.10.2.6. Required Azure permissions for user-provisioned infrastructure

The installation program requires access to an Azure service principal or managed identity with the necessary permissions to deploy the cluster and to maintain its daily operation. These permissions must be granted to the Azure subscription that is associated with the identity.

The following options are available to you:

- You can assign the identity the **Contributor** and **User Access Administrator** roles. Assigning these roles is the quickest way to grant all of the required permissions. For more information about assigning roles, see the Azure documentation for [managing access to Azure resources using the Azure portal](#).
- If your organization's security policies require a more restrictive set of permissions, you can create a [custom role](#) with the necessary permissions.

The following permissions are required for creating an OpenShift Container Platform cluster on Microsoft Azure.

Example 7.37. Required permissions for creating authorization resources

- **Microsoft.Authorization/policies/audit/action**
- **Microsoft.Authorization/policies/auditIfNotExists/action**
- **Microsoft.Authorization/roleAssignments/read**
- **Microsoft.Authorization/roleAssignments/write**

Example 7.38. Required permissions for creating compute resources

- **Microsoft.Compute/images/read**
- **Microsoft.Compute/images/write**
- **Microsoft.Compute/images/delete**
- **Microsoft.Compute/availabilitySets/read**
- **Microsoft.Compute/disks/beginGetAccess/action**
- **Microsoft.Compute/disks/delete**
- **Microsoft.Compute/disks/read**
- **Microsoft.Compute/disks/write**
- **Microsoft.Compute/galleries/images/read**
- **Microsoft.Compute/galleries/images/versions/read**
- **Microsoft.Compute/galleries/images/versions/write**

- **Microsoft.Compute/galleries/images/write**
- **Microsoft.Compute/galleries/read**
- **Microsoft.Compute/galleries/write**
- **Microsoft.Compute/snapshots/read**
- **Microsoft.Compute/snapshots/write**
- **Microsoft.Compute/snapshots/delete**
- **Microsoft.Compute/virtualMachines/delete**
- **Microsoft.Compute/virtualMachines/powerOff/action**
- **Microsoft.Compute/virtualMachines/read**
- **Microsoft.Compute/virtualMachines/write**
- **Microsoft.Compute/virtualMachines/deallocate/action**

Example 7.39. Required permissions for creating identity management resources

- **Microsoft.ManagedIdentity/userAssignedIdentities/assign/action**
- **Microsoft.ManagedIdentity/userAssignedIdentities/read**
- **Microsoft.ManagedIdentity/userAssignedIdentities/write**

Example 7.40. Required permissions for creating network resources

- **Microsoft.Network/dnsZones/A/write**
- **Microsoft.Network/dnsZones/CNAME/write**
- **Microsoft.Network/dnszones/CNAME/read**
- **Microsoft.Network/dnszones/read**
- **Microsoft.Network/loadBalancers/backendAddressPools/join/action**
- **Microsoft.Network/loadBalancers/backendAddressPools/read**
- **Microsoft.Network/loadBalancers/backendAddressPools/write**
- **Microsoft.Network/loadBalancers/read**
- **Microsoft.Network/loadBalancers/write**
- **Microsoft.Network/networkInterfaces/delete**
- **Microsoft.Network/networkInterfaces/join/action**
- **Microsoft.Network/networkInterfaces/read**

- **Microsoft.Network/networkInterfaces/write**
- **Microsoft.Network/networkSecurityGroups/join/action**
- **Microsoft.Network/networkSecurityGroups/read**
- **Microsoft.Network/networkSecurityGroups/securityRules/delete**
- **Microsoft.Network/networkSecurityGroups/securityRules/read**
- **Microsoft.Network/networkSecurityGroups/securityRules/write**
- **Microsoft.Network/networkSecurityGroups/write**
- **Microsoft.Network/privateDnsZones/A/read**
- **Microsoft.Network/privateDnsZones/A/write**
- **Microsoft.Network/privateDnsZones/A/delete**
- **Microsoft.Network/privateDnsZones/SOA/read**
- **Microsoft.Network/privateDnsZones/read**
- **Microsoft.Network/privateDnsZones/virtualNetworkLinks/read**
- **Microsoft.Network/privateDnsZones/virtualNetworkLinks/write**
- **Microsoft.Network/privateDnsZones/write**
- **Microsoft.Network/publicIPAddresses/delete**
- **Microsoft.Network/publicIPAddresses/join/action**
- **Microsoft.Network/publicIPAddresses/read**
- **Microsoft.Network/publicIPAddresses/write**
- **Microsoft.Network/virtualNetworks/join/action**
- **Microsoft.Network/virtualNetworks/read**
- **Microsoft.Network/virtualNetworks/subnets/join/action**
- **Microsoft.Network/virtualNetworks/subnets/read**
- **Microsoft.Network/virtualNetworks/subnets/write**
- **Microsoft.Network/virtualNetworks/write**

Example 7.41. Required permissions for checking the health of resources

- **Microsoft.Resourcehealth/healthevent/Activated/action**
- **Microsoft.Resourcehealth/healthevent/InProgress/action**
- **Microsoft.Resourcehealth/healthevent/Pending/action**

- **Microsoft.Resourcehealth/healthevent/Resolved/action**
- **Microsoft.Resourcehealth/healthevent/Updated/action**

Example 7.42. Required permissions for creating a resource group

- **Microsoft.Resources/subscriptions/resourceGroups/read**
- **Microsoft.Resources/subscriptions/resourcegroups/write**

Example 7.43. Required permissions for creating resource tags

- **Microsoft.Resources/tags/write**

Example 7.44. Required permissions for creating storage resources

- **Microsoft.Storage/storageAccounts/blobServices/read**
- **Microsoft.Storage/storageAccounts/blobServices/containers/write**
- **Microsoft.Storage/storageAccounts/fileServices/read**
- **Microsoft.Storage/storageAccounts/fileServices/shares/read**
- **Microsoft.Storage/storageAccounts/fileServices/shares/write**
- **Microsoft.Storage/storageAccounts/fileServices/shares/delete**
- **Microsoft.Storage/storageAccounts/listKeys/action**
- **Microsoft.Storage/storageAccounts/read**
- **Microsoft.Storage/storageAccounts/write**

Example 7.45. Required permissions for creating deployments

- **Microsoft.Resources/deployments/read**
- **Microsoft.Resources/deployments/write**
- **Microsoft.Resources/deployments/validate/action**
- **Microsoft.Resources/deployments/operationstatuses/read**

Example 7.46. Optional permissions for creating compute resources

- **Microsoft.Compute/availabilitySets/delete**
- **Microsoft.Compute/availabilitySets/write**

Example 7.47. Optional permissions for creating marketplace virtual machine resources

- **Microsoft.MarketplaceOrdering/offertypes/publishers/offers/plans/agreements/read**
- **Microsoft.MarketplaceOrdering/offertypes/publishers/offers/plans/agreements/write**

Example 7.48. Optional permissions for enabling user-managed encryption

- **Microsoft.Compute/diskEncryptionSets/read**
- **Microsoft.Compute/diskEncryptionSets/write**
- **Microsoft.Compute/diskEncryptionSets/delete**
- **Microsoft.KeyVault/vaults/read**
- **Microsoft.KeyVault/vaults/write**
- **Microsoft.KeyVault/vaults/delete**
- **Microsoft.KeyVault/vaults/deploy/action**
- **Microsoft.KeyVault/vaults/keys/read**
- **Microsoft.KeyVault/vaults/keys/write**
- **Microsoft.Features/providers/features/register/action**

The following permissions are required for deleting an OpenShift Container Platform cluster on Microsoft Azure.

Example 7.49. Required permissions for deleting authorization resources

- **Microsoft.Authorization/roleAssignments/delete**

Example 7.50. Required permissions for deleting compute resources

- **Microsoft.Compute/disks/delete**
- **Microsoft.Compute/galleries/delete**
- **Microsoft.Compute/galleries/images/delete**
- **Microsoft.Compute/galleries/images/versions/delete**
- **Microsoft.Compute/virtualMachines/delete**
- **Microsoft.Compute/images/delete**

Example 7.51. Required permissions for deleting identity management resources

- **Microsoft.ManagedIdentity/userAssignedIdentities/delete**

Example 7.52. Required permissions for deleting network resources

- **Microsoft.Network/dnszones/read**
- **Microsoft.Network/dnsZones/A/read**
- **Microsoft.Network/dnsZones/A/delete**
- **Microsoft.Network/dnsZones/CNAME/read**
- **Microsoft.Network/dnsZones/CNAME/delete**
- **Microsoft.Network/loadBalancers/delete**
- **Microsoft.Network/networkInterfaces/delete**
- **Microsoft.Network/networkSecurityGroups/delete**
- **Microsoft.Network/privateDnsZones/read**
- **Microsoft.Network/privateDnsZones/A/read**
- **Microsoft.Network/privateDnsZones/delete**
- **Microsoft.Network/privateDnsZones/virtualNetworkLinks/delete**
- **Microsoft.Network/publicIPAddresses/delete**
- **Microsoft.Network/virtualNetworks/delete**

Example 7.53. Required permissions for checking the health of resources

- **Microsoft.Resourcehealth/healthevent/Activated/action**
- **Microsoft.Resourcehealth/healthevent/Resolved/action**
- **Microsoft.Resourcehealth/healthevent/Updated/action**

Example 7.54. Required permissions for deleting a resource group

- **Microsoft.Resources/subscriptions/resourcegroups/delete**

Example 7.55. Required permissions for deleting storage resources

- **Microsoft.Storage/storageAccounts/delete**
- **Microsoft.Storage/storageAccounts/listKeys/action**



NOTE

To install OpenShift Container Platform on Azure, you must scope the permissions related to resource group creation to your subscription. After the resource group is created, you can scope the rest of the permissions to the created resource group. If the public DNS zone is present in a different resource group, then the network DNS zone related permissions must always be applied to your subscription.

You can scope all the permissions to your subscription when deleting an OpenShift Container Platform cluster.

7.10.2.7. Creating a service principal

Because OpenShift Container Platform and its installation program create Microsoft Azure resources by using the Azure Resource Manager, you must create a service principal to represent it.

Prerequisites

- Install or update the [Azure CLI](#).
- Your Azure account has the required roles for the subscription that you use.
- If you want to use a custom role, you have created a [custom role](#) with the required permissions listed in the *Required Azure permissions for user-provisioned infrastructure* section.

Procedure

1. Log in to the Azure CLI:

```
$ az login
```

2. If your Azure account uses subscriptions, ensure that you are using the right subscription:

- a. View the list of available accounts and record the **tenantId** value for the subscription you want to use for your cluster:

```
$ az account list --refresh
```

Example output

```
[
  {
    "cloudName": "AzureCloud",
    "id": "9bab1460-96d5-40b3-a78e-17b15e978a80",
    "isDefault": true,
    "name": "Subscription Name",
    "state": "Enabled",
    "tenantId": "6057c7e9-b3ae-489d-a54e-de3f6bf6a8ee",
    "user": {
      "name": "you@example.com",
      "type": "user"
    }
  }
]
```

- b. View your active account details and confirm that the **tenantId** value matches the subscription you want to use:

```
$ az account show
```

Example output

```
{
  "environmentName": "AzureCloud",
  "id": "9bab1460-96d5-40b3-a78e-17b15e978a80",
  "isDefault": true,
  "name": "Subscription Name",
  "state": "Enabled",
  "tenantId": "6057c7e9-b3ae-489d-a54e-de3f6bf6a8ee", 1
  "user": {
    "name": "you@example.com",
    "type": "user"
  }
}
```

- 1** Ensure that the value of the **tenantId** parameter is the correct subscription ID.

- c. If you are not using the right subscription, change the active subscription:

```
$ az account set -s <subscription_id> 1
```

- 1** Specify the subscription ID.

- d. Verify the subscription ID update:

```
$ az account show
```

Example output

```
{
  "environmentName": "AzureCloud",
  "id": "33212d16-bdf6-45cb-b038-f6565b61edda",
  "isDefault": true,
  "name": "Subscription Name",
  "state": "Enabled",
  "tenantId": "8049c7e9-c3de-762d-a54e-dc3f6be6a7ee",
  "user": {
    "name": "you@example.com",
    "type": "user"
  }
}
```

- Record the **tenantId** and **id** parameter values from the output. You need these values during the OpenShift Container Platform installation.
- Create the service principal for your account:

```
$ az ad sp create-for-rbac --role <role_name> \ 1
--name <service_principal> \ 2
--scopes /subscriptions/<subscription_id> 3
```

- 1 Defines the role name. You can use the **Contributor** role, or you can specify a custom role which contains the necessary permissions.
- 2 Defines the service principal name.
- 3 Specifies the subscription ID.

Example output

```
Creating 'Contributor' role assignment under scope '/subscriptions/<subscription_id>'
The output includes credentials that you must protect. Be sure that you do not
include these credentials in your code or check the credentials into your source
control. For more information, see https://aka.ms/azadsp-cli
{
  "appId": "ac461d78-bf4b-4387-ad16-7e32e328aec6",
  "displayName": <service_principal>,
  "password": "00000000-0000-0000-0000-000000000000",
  "tenantId": "8049c7e9-c3de-762d-a54e-dc3f6be6a7ee"
}
```

5. Record the values of the **appId** and **password** parameters from the previous output. You need these values during OpenShift Container Platform installation.
6. If you applied the **Contributor** role to your service principal, assign the **User Administrator Access** role by running the following command:

```
$ az role assignment create --role "User Access Administrator" \
--assignee-object-id $(az ad sp show --id <appId> --query id -o tsv) 1
```

- 1 Specify the **appId** parameter value for your service principal.

Additional resources

- For more information about CCO modes, see [About the Cloud Credential Operator](#).

7.10.2.8. Supported Azure regions

The installation program dynamically generates the list of available Microsoft Azure regions based on your subscription.

Supported Azure public regions

- **australiacentral** (Australia Central)
- **australiaeast** (Australia East)
- **australiasoutheast** (Australia South East)
- **brazilsouth** (Brazil South)

- **canadacentral** (Canada Central)
- **canadaeast** (Canada East)
- **centralindia** (Central India)
- **centralus** (Central US)
- **eastasia** (East Asia)
- **eastus** (East US)
- **eastus2** (East US 2)
- **francecentral** (France Central)
- **germanywestcentral** (Germany West Central)
- **israelcentral** (Israel Central)
- **italynorth** (Italy North)
- **japaneast** (Japan East)
- **japanwest** (Japan West)
- **koreacentral** (Korea Central)
- **koreasouth** (Korea South)
- **mexicocentral** (Mexico Central)
- **northcentralus** (North Central US)
- **northeurope** (North Europe)
- **norwayeast** (Norway East)
- **polandcentral** (Poland Central)
- **qatarcentral** (Qatar Central)
- **southafricanorth** (South Africa North)
- **southcentralus** (South Central US)
- **southeastasia** (Southeast Asia)
- **southindia** (South India)
- **swedencentral** (Sweden Central)
- **switzerlandnorth** (Switzerland North)
- **uaenorth** (UAE North)
- **uksouth** (UK South)

- **ukwest** (UK West)
- **westcentralus** (West Central US)
- **westeurope** (West Europe)
- **westindia** (West India)
- **westus** (West US)
- **westus2** (West US 2)
- **westus3** (West US 3)

Supported Azure Government regions

Support for the following Microsoft Azure Government (MAG) regions was added in OpenShift Container Platform version 4.6:

- **usgovtexas** (US Gov Texas)
- **usgovvirginia** (US Gov Virginia)

You can reference all available MAG regions in the [Azure documentation](#). Other provided MAG regions are expected to work with OpenShift Container Platform, but have not been tested.

7.10.3. Requirements for a cluster with user-provisioned infrastructure

For a cluster that contains user-provisioned infrastructure, you must deploy all of the required machines.

This section describes the requirements for deploying OpenShift Container Platform on user-provisioned infrastructure.

7.10.3.1. Required machines for cluster installation

The smallest OpenShift Container Platform clusters require the following hosts:

Table 7.20. Minimum required hosts

Hosts	Description
One temporary bootstrap machine	The cluster requires the bootstrap machine to deploy the OpenShift Container Platform cluster on the three control plane machines. You can remove the bootstrap machine after you install the cluster.
Three control plane machines	The control plane machines run the Kubernetes and OpenShift Container Platform services that form the control plane.
At least two compute machines, which are also known as worker machines.	The workloads requested by OpenShift Container Platform users run on the compute machines.

**IMPORTANT**

To maintain high availability of your cluster, use separate physical hosts for these cluster machines.

The bootstrap and control plane machines must use Red Hat Enterprise Linux CoreOS (RHCOS) as the operating system. However, the compute machines can choose between Red Hat Enterprise Linux CoreOS (RHCOS), Red Hat Enterprise Linux (RHEL) 8.6 and later.

Note that RHCOS is based on Red Hat Enterprise Linux (RHEL) 9.2 and inherits all of its hardware certifications and requirements. See [Red Hat Enterprise Linux technology capabilities and limits](#).

7.10.3.2. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 7.21. Minimum resource requirements

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Bootstrap	RHCOS	4	16 GB	100 GB	300
Control plane	RHCOS	4	16 GB	100 GB	300
Compute	RHCOS, RHEL 8.6 and later [3]	2	8 GB	100 GB	300

1. One vCPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio: $(\text{threads per core} \times \text{cores}) \times \text{sockets} = \text{vCPUs}$.
2. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes which require a 10 ms p99 fsync duration. Note that on many cloud platforms, storage size and IOPS scale together, so you might need to over-allocate storage volume to obtain sufficient performance.
3. As with all user-provisioned installations, if you choose to use RHEL compute machines in your cluster, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Use of RHEL 7 compute machines is deprecated and has been removed in OpenShift Container Platform 4.10 and later.




NOTE

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).



IMPORTANT

You are required to use Azure virtual machines that have the **premiumIO** parameter set to **true**.

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

7.10.3.3. Tested instance types for Azure

The following Microsoft Azure instance types have been tested with OpenShift Container Platform.

Example 7.56. Machine types based on 64-bit x86 architecture

- **c4.***
- **c5.***
- **c5a.***
- **i3.***
- **m4.***
- **m5.***
- **m5a.***
- **m6a.***
- **m6i.***
- **r4.***
- **r5.***
- **r5a.***
- **r6i.***

- **t3.***
- **t3a.***

7.10.3.4. Tested instance types for Azure on 64-bit ARM infrastructures

The following Microsoft Azure ARM64 instance types have been tested with OpenShift Container Platform.

Example 7.57. Machine types based on 64-bit ARM architecture

- **c6g.***
- **m6g.***

7.10.4. Using the Azure Marketplace offering

Using the Azure Marketplace offering lets you deploy an OpenShift Container Platform cluster, which is billed on pay-per-use basis (hourly, per core) through Azure, while still being supported directly by Red Hat.

To deploy an OpenShift Container Platform cluster using the Azure Marketplace offering, you must first obtain the Azure Marketplace image. The installation program uses this image to deploy worker or control plane nodes. When obtaining your image, consider the following:

- While the images are the same, the Azure Marketplace publisher is different depending on your region. If you are located in North America, specify **redhat** as the publisher. If you are located in EMEA, specify **redhat-limited** as the publisher.
- The offer includes a **rh-ocp-worker** SKU and a **rh-ocp-worker-gen1** SKU. The **rh-ocp-worker** SKU represents a Hyper-V generation version 2 VM image. The default instance types used in OpenShift Container Platform are version 2 compatible. If you plan to use an instance type that is only version 1 compatible, use the image associated with the **rh-ocp-worker-gen1** SKU. The **rh-ocp-worker-gen1** SKU represents a Hyper-V version 1 VM image.



IMPORTANT

Installing images with the Azure marketplace is not supported on clusters with 64-bit ARM instances.

Prerequisites

- You have installed the Azure CLI client (**az**).
- Your Azure account is entitled for the offer and you have logged into this account with the Azure CLI client.

Procedure

1. Display all of the available OpenShift Container Platform images by running one of the following commands:
 - North America:

```
$ az vm image list --all --offer rh-ocp-worker --publisher redhat -o table
```

Example output

Offer	Publisher	Sku	Urn	Version
rh-ocp-worker	RedHat	rh-ocp-worker	RedHat:rh-ocp-worker:rh-ocp-worker:413.92.2023101700	413.92.2023101700
rh-ocp-worker-gen1	RedHat	rh-ocp-worker-gen1	RedHat:rh-ocp-worker:rh-ocp-worker-gen1:413.92.2023101700	413.92.2023101700

- EMEA:

```
$ az vm image list --all --offer rh-ocp-worker --publisher redhat-limited -o table
```

Example output

Offer	Publisher	Sku	Urn	Version
rh-ocp-worker	redhat-limited	rh-ocp-worker	redhat-limited:rh-ocp-worker:rh-ocp-worker:413.92.2023101700	413.92.2023101700
rh-ocp-worker-gen1	redhat-limited	rh-ocp-worker-gen1	redhat-limited:rh-ocp-worker:rh-ocp-worker-gen1:413.92.2023101700	413.92.2023101700



NOTE

Use the latest image that is available for compute and control plane nodes. If required, your VMs are automatically upgraded as part of the installation process.

2. Inspect the image for your offer by running one of the following commands:

- North America:

```
$ az vm image show --urn redhat:rh-ocp-worker:rh-ocp-worker:<version>
```

- EMEA:

```
$ az vm image show --urn redhat-limited:rh-ocp-worker:rh-ocp-worker:<version>
```

3. Review the terms of the offer by running one of the following commands:

- North America:

```
$ az vm image terms show --urn redhat:rh-ocp-worker:rh-ocp-worker:<version>
```

- EMEA:

```
$ az vm image terms show --urn redhat-limited:rh-ocp-worker:rh-ocp-worker:<version>
```

4. Accept the terms of the offering by running one of the following commands:

- North America:

```
$ az vm image terms accept --urn redhat:rh-ocp-worker:rh-ocp-worker:<version>
```

- EMEA:

```
$ az vm image terms accept --urn redhat-limited:rh-ocp-worker:rh-ocp-worker:<version>
```

5. Record the image details of your offer. If you use the Azure Resource Manager (ARM) template to deploy your compute nodes:

- Update **storageProfile.imageReference** by deleting the **id** parameter and adding the **offer**, **publisher**, **sku**, and **version** parameters by using the values from your offer.
- Specify a **plan** for the virtual machines (VMs).

Example 06_workers.json ARM template with an updated storageProfile.imageReference object and a specified plan

```
...
  "plan" : {
    "name": "rh-ocp-worker",
    "product": "rh-ocp-worker",
    "publisher": "redhat"
  },
  "dependsOn" : [
    "[concat('Microsoft.Network/networkInterfaces/', concat(variables('vmNames')
[copyIndex()], '-nic'))]"
  ],
  "properties" : {
...
  "storageProfile": {
    "imageReference": {
      "offer": "rh-ocp-worker",
      "publisher": "redhat",
      "sku": "rh-ocp-worker",
      "version": "413.92.2023101700"
    }
...
  }
...
}
```

7.10.4.1. Obtaining the installation program

Before you install OpenShift Container Platform, download the installation file on the host you are using for installation.

Prerequisites

- You have a computer that runs Linux or macOS, with 500 MB of local disk space.


Procedure

Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select your infrastructure provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.

**IMPORTANT**

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.

**IMPORTANT**

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.

4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#). This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

7.10.4.2. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.

**IMPORTANT**

Do not skip this procedure in production environments, where disaster recovery and debugging is required.

**NOTE**

You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.

**NOTE**

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.

**NOTE**

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```

**NOTE**

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1 Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program. If you install a cluster on infrastructure that you provision, you must provide the key to the installation program.

7.10.5. Creating the installation files for Azure

To install OpenShift Container Platform on Microsoft Azure using user-provisioned infrastructure, you must generate the files that the installation program needs to deploy your cluster and modify them so that the cluster creates only the machines that it will use. You generate and customize the **install-config.yaml** file, Kubernetes manifests, and Ignition config files. You also have the option to first set up a separate **var** partition during the preparation phases of installation.

7.10.5.1. Optional: Creating a separate /var partition

It is recommended that disk partitioning for OpenShift Container Platform be left to the installer. However, there are cases where you might want to create separate partitions in a part of the filesystem that you expect to grow.

OpenShift Container Platform supports the addition of a single partition to attach storage to either the **/var** partition or a subdirectory of **/var**. For example:

- **/var/lib/containers**: Holds container-related content that can grow as more images and containers are added to a system.
- **/var/lib/etcd**: Holds data that you might want to keep separate for purposes such as performance optimization of etcd storage.
- **/var**: Holds data that you might want to keep separate for purposes such as auditing.

Storing the contents of a **/var** directory separately makes it easier to grow storage for those areas as needed and reinstall OpenShift Container Platform at a later date and keep that data intact. With this method, you will not have to pull all your containers again, nor will you have to copy massive log files when you update systems.

Because **/var** must be in place before a fresh installation of Red Hat Enterprise Linux CoreOS (RHCOS), the following procedure sets up the separate **/var** partition by creating a machine config manifest that is

inserted during the **openshift-install** preparation phases of an OpenShift Container Platform installation.



IMPORTANT

If you follow the steps to create a separate **/var** partition in this procedure, it is not necessary to create the Kubernetes manifest and Ignition config files again as described later in this section.

Procedure

1. Create a directory to hold the OpenShift Container Platform installation files:

```
$ mkdir $HOME/clusterconfig
```

2. Run **openshift-install** to create a set of files in the **manifest** and **openshift** subdirectories. Answer the system questions as you are prompted:

```
$ openshift-install create manifests --dir $HOME/clusterconfig
```

Example output

```
? SSH Public Key ...
INFO Credentials loaded from the "myprofile" profile in file "/home/myuser/.aws/credentials"
INFO Consuming Install Config from target directory
INFO Manifests created in: $HOME/clusterconfig/manifests and
$HOME/clusterconfig/openshift
```

3. Optional: Confirm that the installation program created manifests in the **clusterconfig/openshift** directory:

```
$ ls $HOME/clusterconfig/openshift/
```

Example output

```
99_kubeadmin-password-secret.yaml
99_openshift-cluster-api_master-machines-0.yaml
99_openshift-cluster-api_master-machines-1.yaml
99_openshift-cluster-api_master-machines-2.yaml
...
```

4. Create a Butane config that configures the additional partition. For example, name the file **\$HOME/clusterconfig/98-var-partition.bu**, change the disk device name to the name of the storage device on the **worker** systems, and set the storage size as appropriate. This example places the **/var** directory on a separate partition:

```
variant: openshift
version: 4.16.0
metadata:
  labels:
    machineconfiguration.openshift.io/role: worker
  name: 98-var-partition
storage:
```



```

disks:
- device: /dev/disk/by-id/<device_name> ❶
  partitions:
  - label: var
    start_mib: <partition_start_offset> ❷
    size_mib: <partition_size> ❸
    number: 5
  filesystems:
  - device: /dev/disk/by-partlabel/var
    path: /var
    format: xfs
    mount_options: [defaults, prjquota] ❹
    with_mount_unit: true

```

- ❶ The storage device name of the disk that you want to partition.
- ❷ When adding a data partition to the boot disk, a minimum value of 25000 MiB (Mebibytes) is recommended. The root file system is automatically resized to fill all available space up to the specified offset. If no value is specified, or if the specified value is smaller than the recommended minimum, the resulting root file system will be too small, and future reinstalls of RHCOS might overwrite the beginning of the data partition.
- ❸ The size of the data partition in mebibytes.
- ❹ The **prjquota** mount option must be enabled for filesystems used for container storage.



NOTE

When creating a separate **/var** partition, you cannot use different instance types for worker nodes, if the different instance types do not have the same device name.

5. Create a manifest from the Butane config and save it to the **clusterconfig/openshift** directory. For example, run the following command:

```
$ butane $HOME/clusterconfig/98-var-partition.bu -o $HOME/clusterconfig/openshift/98-var-partition.yaml
```

6. Run **openshift-install** again to create Ignition configs from a set of files in the **manifest** and **openshift** subdirectories:

```
$ openshift-install create ignition-configs --dir $HOME/clusterconfig
$ ls $HOME/clusterconfig/
auth bootstrap.ign master.ign metadata.json worker.ign
```

Now you can use the Ignition config files as input to the installation procedures to install Red Hat Enterprise Linux CoreOS (RHCOS) systems.

7.10.5.2. Creating the installation configuration file

You can customize the OpenShift Container Platform cluster you install on Microsoft Azure.

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster. For a restricted network installation, these files are on your mirror host.
- You have the **imageContentSources** values that were generated during mirror registry creation.
- You have obtained the contents of the certificate for your mirror registry.
- You have retrieved a Red Hat Enterprise Linux CoreOS (RHCOS) image and uploaded it to an accessible location.
- You have an Azure subscription ID and tenant ID.
- If you are installing the cluster using a service principal, you have its application ID and password.
- If you are installing the cluster using a system-assigned managed identity, you have enabled it on the virtual machine that you will run the installation program from.
- If you are installing the cluster using a user-assigned managed identity, you have met these prerequisites:
 - You have its client ID.
 - You have assigned it to the virtual machine that you will run the installation program from.

Procedure

1. Optional: If you have run the installation program on this computer before, and want to use an alternative service principal or managed identity, go to the `~/.azure/` directory and delete the **osServicePrincipal.json** configuration file.

Deleting this file prevents the installation program from automatically reusing subscription and authentication values from a previous installation.

2. Create the **install-config.yaml** file.
 - a. Change to the directory that contains the installation program and run the following command:

```
┆ $ ./openshift-install create install-config --dir <installation_directory> 1
```

- 1 For **<installation_directory>**, specify the directory name to store the files that the installation program creates.

When specifying the directory:

- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
- Use an empty directory. Some installation assets, such as bootstrap X.509 certificates, have short expiration intervals, therefore you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

- b. At the prompts, provide the configuration details for your cloud:
- i. Optional: Select an SSH key to use to access your cluster machines.

**NOTE**

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- ii. Select **azure** as the platform to target.
If the installation program cannot locate the **osServicePrincipal.json** configuration file from a previous installation, you are prompted for Azure subscription and authentication values.
- iii. Enter the following Azure parameter values for your subscription:
 - **azure subscription id** Enter the subscription ID to use for the cluster.
 - **azure tenant id** Enter the tenant ID.
- iv. Depending on the Azure identity you are using to deploy the cluster, do one of the following when prompted for the **azure service principal client id**
 - If you are using a service principal, enter its application ID.
 - If you are using a system-assigned managed identity, leave this value blank.
 - If you are using a user-assigned managed identity, specify its client ID.
- v. Depending on the Azure identity you are using to deploy the cluster, do one of the following when prompted for the **azure service principal client secret**
 - If you are using a service principal, enter its password.
 - If you are using a system-assigned managed identity, leave this value blank.
 - If you are using a user-assigned managed identity, leave this value blank.
- vi. Select the region to deploy the cluster to.
- vii. Select the base domain to deploy the cluster to. The base domain corresponds to the Azure DNS Zone that you created for your cluster.
- viii. Enter a descriptive name for your cluster.

**IMPORTANT**

All Azure resources that are available through public endpoints are subject to resource name restrictions, and you cannot create resources that use certain terms. For a list of terms that Azure restricts, see [Resolve reserved resource name errors](#) in the Azure documentation.

- ix. Paste the [pull secret from Red Hat OpenShift Cluster Manager](#) .

3. Edit the **install-config.yaml** file to give the additional information that is required for an installation in a restricted network.

- a. Update the **pullSecret** value to contain the authentication information for your registry:

```
pullSecret: '{"auths":{"<mirror_host_name>:5000": {"auth": "<credentials>","email":
"you@example.com"}}}'
```

For **<mirror_host_name>**, specify the registry domain name that you specified in the certificate for your mirror registry, and for **<credentials>**, specify the base64-encoded user name and password for your mirror registry.

- b. Add the **additionalTrustBundle** parameter and value.

```
additionalTrustBundle: |
  -----BEGIN CERTIFICATE-----
  ////////////////////////////////////////////////////////////////////
  -----END CERTIFICATE-----
```

The value must be the contents of the certificate file that you used for your mirror registry. The certificate file can be an existing, trusted certificate authority, or the self-signed certificate that you generated for the mirror registry.

- c. Define the network and subnets for the VNet to install the cluster under the **platform.azure** field:

```
networkResourceGroupName: <vnet_resource_group> 1
virtualNetwork: <vnet> 2
controlPlaneSubnet: <control_plane_subnet> 3
computeSubnet: <compute_subnet> 4
```

- 1** Replace **<vnet_resource_group>** with the resource group name that contains the existing virtual network (VNet).
- 2** Replace **<vnet>** with the existing virtual network name.
- 3** Replace **<control_plane_subnet>** with the existing subnet name to deploy the control plane machines.
- 4** Replace **<compute_subnet>** with the existing subnet name to deploy compute machines.

- d. Add the image content resources, which resemble the following YAML excerpt:

```
imageContentSources:
- mirrors:
  - <mirror_host_name>:5000/<repo_name>/release
    source: quay.io/openshift-release-dev/ocp-release
- mirrors:
  - <mirror_host_name>:5000/<repo_name>/release
    source: registry.redhat.io/ocp/release
```

For these values, use the **imageContentSources** that you recorded during mirror registry creation.

- e. Optional: Set the publishing strategy to **Internal**:

```
publish: Internal
```

By setting this option, you create an internal Ingress Controller and a private load balancer.



IMPORTANT

Azure Firewall [does not work seamlessly](#) with Azure Public Load balancers. Thus, when using Azure Firewall for restricting internet access, the **publish** field in **install-config.yaml** should be set to **Internal**.

4. Make any other modifications to the **install-config.yaml** file that you require. For more information about the parameters, see "Installation configuration parameters".
5. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

If previously not detected, the installation program creates an **osServicePrincipal.json** configuration file and stores this file in the `~/.azure/` directory on your computer. This ensures that the installation program can load the profile when it is creating an OpenShift Container Platform cluster on the target platform.

7.10.5.3. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:


```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.



NOTE

The installation program does not support the proxy **readinessEndpoints** field.

**NOTE**

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.

**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

7.10.5.4. Exporting common variables for ARM templates

You must export a common set of variables that are used with the provided Azure Resource Manager (ARM) templates used to assist in completing a user-provided infrastructure install on Microsoft Azure.

**NOTE**

Specific ARM templates can also require additional exported variables, which are detailed in their related procedures.

Prerequisites

- Obtain the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Export common variables found in the **install-config.yaml** to be used by the provided ARM templates:

```
$ export CLUSTER_NAME=<cluster_name> 1
$ export AZURE_REGION=<azure_region> 2
$ export SSH_KEY=<ssh_key> 3
$ export BASE_DOMAIN=<base_domain> 4
$ export BASE_DOMAIN_RESOURCE_GROUP=<base_domain_resource_group> 5
```

- 1 The value of the **.metadata.name** attribute from the **install-config.yaml** file.
- 2 The region to deploy the cluster into, for example **centralus**. This is the value of the **.platform.azure.region** attribute from the **install-config.yaml** file.
- 3 The SSH RSA public key file as a string. You must enclose the SSH key in quotes since it contains spaces. This is the value of the **.sshKey** attribute from the **install-config.yaml** file.

- 4 The base domain to deploy the cluster to. The base domain corresponds to the public DNS zone that you created for your cluster. This is the value of the **.baseDomain** attribute
- 5 The resource group where the public DNS zone exists. This is the value of the **.platform.azure.baseDomainResourceGroupName** attribute from the **install-config.yaml** file.

For example:

```
$ export CLUSTER_NAME=test-cluster
$ export AZURE_REGION=centralus
$ export SSH_KEY="ssh-rsa xxx/xxx/xxx= user@email.com"
$ export BASE_DOMAIN=example.com
$ export BASE_DOMAIN_RESOURCE_GROUP=ocp-cluster
```

2. Export the kubeadmin credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

7.10.5.5. Creating the Kubernetes manifest and Ignition config files

Because you must modify some cluster definition files and manually start the cluster machines, you must generate the Kubernetes manifest and Ignition config files that the cluster needs to configure the machines.

The installation configuration file transforms into the Kubernetes manifests. The manifests wrap into the Ignition configuration files, which are later used to configure the cluster machines.



IMPORTANT

- The Ignition config files that the OpenShift Container Platform installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

Prerequisites

- You obtained the OpenShift Container Platform installation program.
- You created the **install-config.yaml** installation configuration file.

Procedure

1. Change to the directory that contains the OpenShift Container Platform installation program and generate the Kubernetes manifests for the cluster:

```
$ ./openshift-install create manifests --dir <installation_directory> 1
```

- 1 For **<installation_directory>**, specify the installation directory that contains the **install-config.yaml** file you created.

2. Remove the Kubernetes manifest files that define the control plane machines:

```
$ rm -f <installation_directory>/openshift/99_openshift-cluster-api_master-machines-*.yaml
```


By removing these files, you prevent the cluster from automatically generating control plane machines.

3. Remove the Kubernetes manifest files that define the control plane machine set:

```
$ rm -f <installation_directory>/openshift/99_openshift-machine-api_master-control-plane-machine-set.yaml
```

4. Remove the Kubernetes manifest files that define the worker machines:

```
$ rm -f <installation_directory>/openshift/99_openshift-cluster-api_worker-machineset-*.yaml
```



IMPORTANT

If you disabled the **MachineAPI** capability when installing a cluster on user-provisioned infrastructure, you must remove the Kubernetes manifest files that define the worker machines. Otherwise, your cluster fails to install.

Because you create and manage the worker machines yourself, you do not need to initialize these machines.

5. Check that the **mastersSchedulable** parameter in the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** Kubernetes manifest file is set to **false**. This setting prevents pods from being scheduled on the control plane machines:
 - a. Open the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** file.
 - b. Locate the **mastersSchedulable** parameter and ensure that it is set to **false**.
 - c. Save and exit the file.
6. Optional: If you do not want [the Ingress Operator](#) to create DNS records on your behalf, remove the **privateZone** and **publicZone** sections from the **<installation_directory>/manifests/cluster-dns-02-config.yml** DNS configuration file:

```
apiVersion: config.openshift.io/v1
kind: DNS
metadata:
```

```

creationTimestamp: null
name: cluster
spec:
  baseDomain: example.openshift.com
  privateZone: ❶
    id: mycluster-100419-private-zone
  publicZone: ❷
    id: example.openshift.com
status: {}

```

- ❶ ❷ Remove this section completely.

If you do so, you must add ingress DNS records manually in a later step.

7. When configuring Azure on user-provisioned infrastructure, you must export some common variables defined in the manifest files to use later in the Azure Resource Manager (ARM) templates:

- a. Export the infrastructure ID by using the following command:

```
$ export INFRA_ID=<infra_id> ❶
```

- ❶ The OpenShift Container Platform cluster has been assigned an identifier (**INFRA_ID**) in the form of **<cluster_name>-<random_string>**. This will be used as the base name for most resources created using the provided ARM templates. This is the value of the **.status.infrastructureName** attribute from the **manifests/cluster-infrastructure-02-config.yml** file.

- b. Export the resource group by using the following command:

```
$ export RESOURCE_GROUP=<resource_group> ❶
```

- ❶ All resources created in this Azure deployment exists as part of a [resource group](#). The resource group name is also based on the **INFRA_ID**, in the form of **<cluster_name>-<random_string>-rg**. This is the value of the **.status.platformStatus.azure.resourceGroupName** attribute from the **manifests/cluster-infrastructure-02-config.yml** file.

8. To create the Ignition configuration files, run the following command from the directory that contains the installation program:

```
$ ./openshift-install create ignition-configs --dir <installation_directory> ❶
```


- ❶ For **<installation_directory>**, specify the same installation directory.

Ignition config files are created for the bootstrap, control plane, and compute nodes in the installation directory. The **kubeadmin-password** and **kubeconfig** files are created in the **./<installation_directory>/auth** directory:

```

.
├── auth

```



7.10.6. Creating the Azure resource group

You must create a Microsoft Azure [resource group](#) and an identity for that resource group. These are both used during the installation of your OpenShift Container Platform cluster on Azure.

Prerequisites

- Configure an Azure account.
- Generate the Ignition config files for your cluster.

Procedure

1. Create the resource group in a supported Azure region:

```
$ az group create --name ${RESOURCE_GROUP} --location ${AZURE_REGION}
```

2. Create an Azure identity for the resource group:

```
$ az identity create -g ${RESOURCE_GROUP} -n ${INFRA_ID}-identity
```

This is used to grant the required access to Operators in your cluster. For example, this allows the Ingress Operator to create a public IP and its load balancer. You must assign the Azure identity to a role.

3. Grant the Contributor role to the Azure identity:

- a. Export the following variables required by the Azure role assignment:

```
$ export PRINCIPAL_ID=`az identity show -g ${RESOURCE_GROUP} -n ${INFRA_ID}-identity --query principalId --out tsv`
```

```
$ export RESOURCE_GROUP_ID=`az group show -g ${RESOURCE_GROUP} --query id --out tsv`
```

- b. Assign the Contributor role to the identity:

```
$ az role assignment create --assignee "${PRINCIPAL_ID}" --role 'Contributor' --scope "${RESOURCE_GROUP_ID}"
```

**NOTE**

If you want to assign a custom role with all the required permissions to the identity, run the following command:

```
$ az role assignment create --assignee "${PRINCIPAL_ID}" --role
<custom_role> \ 1
--scope "${RESOURCE_GROUP_ID}"
```

1 Specifies the custom role name.

7.10.7. Uploading the RHCOS cluster image and bootstrap Ignition config file

The Azure client does not support deployments based on files existing locally. You must copy and store the RHCOS virtual hard disk (VHD) cluster image and bootstrap Ignition config file in a storage container so they are accessible during deployment.

Prerequisites

- Configure an Azure account.
- Generate the Ignition config files for your cluster.

Procedure

1. Create an Azure storage account to store the VHD cluster image:

```
$ az storage account create -g ${RESOURCE_GROUP} --location ${AZURE_REGION} --
name ${CLUSTER_NAME}sa --kind Storage --sku Standard_LRS
```

**WARNING**

The Azure storage account name must be between 3 and 24 characters in length and use numbers and lower-case letters only. If your **CLUSTER_NAME** variable does not follow these restrictions, you must manually define the Azure storage account name. For more information on Azure storage account name restrictions, see [Resolve errors for storage account names](#) in the Azure documentation.

2. Export the storage account key as an environment variable:

```
$ export ACCOUNT_KEY=`az storage account keys list -g ${RESOURCE_GROUP} --
account-name ${CLUSTER_NAME}sa --query "[0].value" -o tsv`
```

3. Export the URL of the RHCOS VHD to an environment variable:

```
$ export VHD_URL=`openshift-install coreos print-stream-json | jq -r '.architectures.
<architecture>."rhel-coreos-extensions"."azure-disk".url`
```

where:

<architecture>

Specifies the architecture, valid values include **x86_64** or **aarch64**.



IMPORTANT

The RHCOS images might not change with every release of OpenShift Container Platform. You must specify an image with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Use the image version that matches your OpenShift Container Platform version if it is available.

4. Create the storage container for the VHD:

```
$ az storage container create --name vhd --account-name ${CLUSTER_NAME}sa --account-key ${ACCOUNT_KEY}
```

5. Copy the local VHD to a blob:

```
$ az storage blob copy start --account-name ${CLUSTER_NAME}sa --account-key ${ACCOUNT_KEY} --destination-blob "rhcos.vhd" --destination-container vhd --source-uri "${VHD_URL}"
```

6. Create a blob storage container and upload the generated **bootstrap.ign** file:


```
$ az storage container create --name files --account-name ${CLUSTER_NAME}sa --account-key ${ACCOUNT_KEY}
```

```
$ az storage blob upload --account-name ${CLUSTER_NAME}sa --account-key ${ACCOUNT_KEY} -c "files" -f "<installation_directory>/bootstrap.ign" -n "bootstrap.ign"
```

7.10.8. Example for creating DNS zones

DNS records are required for clusters that use user-provisioned infrastructure. You should choose the DNS strategy that fits your scenario.

For this example, [Azure's DNS solution](#) is used, so you will create a new public DNS zone for external (internet) visibility and a private DNS zone for internal cluster resolution.



NOTE

The public DNS zone is not required to exist in the same resource group as the cluster deployment and might already exist in your organization for the desired base domain. If that is the case, you can skip creating the public DNS zone; be sure the installation config you generated earlier reflects that scenario.

Prerequisites

- Configure an Azure account.

- Generate the Ignition config files for your cluster.

Procedure

1. Create the new public DNS zone in the resource group exported in the **BASE_DOMAIN_RESOURCE_GROUP** environment variable:

```
$ az network dns zone create -g ${BASE_DOMAIN_RESOURCE_GROUP} -n
${CLUSTER_NAME}.${BASE_DOMAIN}
```

You can skip this step if you are using a public DNS zone that already exists.

2. Create the private DNS zone in the same resource group as the rest of this deployment:

```
$ az network private-dns zone create -g ${RESOURCE_GROUP} -n
${CLUSTER_NAME}.${BASE_DOMAIN}
```

You can learn more about [configuring a public DNS zone in Azure](#) by visiting that section.

7.10.9. Creating a VNet in Azure

You must create a virtual network (VNet) in Microsoft Azure for your OpenShift Container Platform cluster to use. You can customize the VNet to meet your requirements. One way to create the VNet is to modify the provided Azure Resource Manager (ARM) template.



NOTE

If you do not use the provided ARM template to create your Azure infrastructure, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- Configure an Azure account.
- Generate the Ignition config files for your cluster.

Procedure

1. Copy the template from the **ARM template for the VNet** section of this topic and save it as **01_vnet.json** in your cluster's installation directory. This template describes the VNet that your cluster requires.
2. Create the deployment by using the **az** CLI:

```
$ az deployment group create -g ${RESOURCE_GROUP} \
  --template-file "<installation_directory>/01_vnet.json" \
  --parameters baseName="${INFRA_ID}" 1
```

- 1** The base name to be used in resource names; this is usually the cluster's infrastructure ID.

3. Link the VNet template to the private DNS zone:

```
$ az network private-dns link vnet create -g ${RESOURCE_GROUP} -z
${CLUSTER_NAME}.${BASE_DOMAIN} -n ${INFRA_ID}-network-link -v "${INFRA_ID}-vnet"
-e false
```

7.10.9.1. ARM template for the VNet

You can use the following Azure Resource Manager (ARM) template to deploy the VNet that you need for your OpenShift Container Platform cluster:

Example 7.58. 01_vnet.json ARM template

```
{
  "$schema" : "https://schema.management.azure.com/schemas/2015-01-01/deploymentTemplate.json#",
  "contentVersion" : "1.0.0.0",
  "parameters" : {
    "baseName" : {
      "type" : "string",
      "minLength" : 1,
      "metadata" : {
        "description" : "Base name to be used in resource names (usually the cluster's Infra ID)"
      }
    }
  },
  "variables" : {
    "location" : "[resourceGroup().location]",
    "virtualNetworkName" : "[concat(parameters('baseName'), '-vnet')]",
    "addressPrefix" : "10.0.0.0/16",
    "masterSubnetName" : "[concat(parameters('baseName'), '-master-subnet')]",
    "masterSubnetPrefix" : "10.0.0.0/24",
    "nodeSubnetName" : "[concat(parameters('baseName'), '-worker-subnet')]",
    "nodeSubnetPrefix" : "10.0.1.0/24",
    "clusterNsgName" : "[concat(parameters('baseName'), '-nsg')]"
  },
  "resources" : [
    {
      "apiVersion" : "2018-12-01",
      "type" : "Microsoft.Network/virtualNetworks",
      "name" : "[variables('virtualNetworkName')]",
      "location" : "[variables('location')]",
      "dependsOn" : [
        "[concat('Microsoft.Network/networkSecurityGroups/', variables('clusterNsgName'))]"
      ],
      "properties" : {
        "addressSpace" : {
          "addressPrefixes" : [
            "[variables('addressPrefix')]"
          ]
        },
        "subnets" : [
          {
            "name" : "[variables('masterSubnetName')]",
            "properties" : {
              "addressPrefix" : "[variables('masterSubnetPrefix')]",
              "serviceEndpoints": [],
```

```

        "networkSecurityGroup" : {
            "id" : "[resourceId('Microsoft.Network/networkSecurityGroups',
variables('clusterNsgName'))]"
        }
    },
    {
        "name" : "[variables('nodeSubnetName')]",
        "properties" : {
            "addressPrefix" : "[variables('nodeSubnetPrefix')]",
            "serviceEndpoints" : [],
            "networkSecurityGroup" : {
                "id" : "[resourceId('Microsoft.Network/networkSecurityGroups',
variables('clusterNsgName'))]"
            }
        }
    }
]
},
{
    "type" : "Microsoft.Network/networkSecurityGroups",
    "name" : "[variables('clusterNsgName')]",
    "apiVersion" : "2018-10-01",
    "location" : "[variables('location')]",
    "properties" : {
        "securityRules" : [
            {
                "name" : "apiserver_in",
                "properties" : {
                    "protocol" : "Tcp",
                    "sourcePortRange" : "*",
                    "destinationPortRange" : "6443",
                    "sourceAddressPrefix" : "*",
                    "destinationAddressPrefix" : "*",
                    "access" : "Allow",
                    "priority" : 101,
                    "direction" : "Inbound"
                }
            }
        ]
    }
}
]
}
}

```

7.10.10. Deploying the RHCOS cluster image for the Azure infrastructure

You must use a valid Red Hat Enterprise Linux CoreOS (RHCOS) image for Microsoft Azure for your OpenShift Container Platform nodes.

Prerequisites

- Configure an Azure account.

- Generate the Ignition config files for your cluster.
- Store the RHCOS virtual hard disk (VHD) cluster image in an Azure storage container.
- Store the bootstrap Ignition config file in an Azure storage container.

Procedure

1. Copy the template from the **ARM template for image storage** section of this topic and save it as **02_storage.json** in your cluster's installation directory. This template describes the image storage that your cluster requires.
2. Export the RHCOS VHD blob URL as a variable:

```
$ export VHD_BLOB_URL=`az storage blob url --account-name ${CLUSTER_NAME}sa --
account-key ${ACCOUNT_KEY} -c vhd -n "rhcos.vhd" -o tsv`
```

3. Deploy the cluster image:

```
$ az deployment group create -g ${RESOURCE_GROUP} \
--template-file "<installation_directory>/02_storage.json" \
--parameters vhdBlobURL="${VHD_BLOB_URL}" \ ❶
--parameters baseName="${INFRA_ID}" \ ❷
--parameters storageAccount="${CLUSTER_NAME}sa" \ ❸
--parameters architecture="<architecture>" ❹
```

- ❶ The blob URL of the RHCOS VHD to be used to create master and worker machines.
- ❷ The base name to be used in resource names; this is usually the cluster's infrastructure ID.
- ❸ The name of your Azure storage account.
- ❹ Specify the system architecture. Valid values are **x64** (default) or **Arm64**.

7.10.10.1. ARM template for image storage

You can use the following Azure Resource Manager (ARM) template to deploy the stored Red Hat Enterprise Linux CoreOS (RHCOS) image that you need for your OpenShift Container Platform cluster:

Example 7.59. 02_storage.json ARM template

```
{
  "$schema": "https://schema.management.azure.com/schemas/2019-04-
01/deploymentTemplate.json#",
  "contentVersion": "1.0.0.0",
  "parameters": {
    "architecture": {
      "type": "string",
      "metadata": {
        "description": "The architecture of the Virtual Machines"
      },
      "defaultValue": "x64",
      "allowedValues": [
```

```

    "Arm64",
    "x64"
  ]
},
"baseName": {
  "type": "string",
  "minLength": 1,
  "metadata": {
    "description": "Base name to be used in resource names (usually the cluster's Infra ID)"
  }
},
"storageAccount": {
  "type": "string",
  "metadata": {
    "description": "The Storage Account name"
  }
},
"vhdBlobURL": {
  "type": "string",
  "metadata": {
    "description": "URL pointing to the blob where the VHD to be used to create master and
worker machines is located"
  }
}
},
"variables": {
  "location": "[resourceGroup().location]",
  "galleryName": "[concat('gallery_', replace(parameters('baseName'), '-', '_'))]",
  "imageName": "[parameters('baseName')]",
  "imageNameGen2": "[concat(parameters('baseName'), '-gen2')]",
  "imageRelease": "1.0.0"
},
"resources": [
  {
    "apiVersion": "2021-10-01",
    "type": "Microsoft.Compute/galleries",
    "name": "[variables('galleryName')]",
    "location": "[variables('location')]",
    "resources": [
      {
        "apiVersion": "2021-10-01",
        "type": "images",
        "name": "[variables('imageName')]",
        "location": "[variables('location')]",
        "dependsOn": [
          "[variables('galleryName')]"
        ],
        "properties": {
          "architecture": "[parameters('architecture')]",
          "hyperVGeneration": "V1",
          "identifier": {
            "offer": "rhcos",
            "publisher": "RedHat",
            "sku": "basic"
          },
          "osState": "Generalized",

```

```

    "osType": "Linux"
  },
  "resources": [
    {
      "apiVersion": "2021-10-01",
      "type": "versions",
      "name": "[variables('imageRelease')]",
      "location": "[variables('location')]",
      "dependsOn": [
        "[variables('imageName')]"
      ],
      "properties": {
        "publishingProfile": {
          "storageAccountType": "Standard_LRS",
          "targetRegions": [
            {
              "name": "[variables('location')]",
              "regionalReplicaCount": "1"
            }
          ]
        },
        "storageProfile": {
          "osDiskImage": {
            "source": {
              "id": "[resourceId('Microsoft.Storage/storageAccounts',
parameters('storageAccount'))]",
              "uri": "[parameters('vhdBlobURL')]"
            }
          }
        }
      }
    }
  ],
  {
    "apiVersion": "2021-10-01",
    "type": "images",
    "name": "[variables('imageNameGen2')]",
    "location": "[variables('location')]",
    "dependsOn": [
      "[variables('galleryName')]"
    ],
    "properties": {
      "architecture": "[parameters('architecture')]",
      "hyperVGeneration": "V2",
      "identifier": {
        "offer": "rhcos-gen2",
        "publisher": "RedHat-gen2",
        "sku": "gen2"
      },
      "osState": "Generalized",
      "osType": "Linux"
    }
  },
  "resources": [
    {
      "apiVersion": "2021-10-01",

```

```
"type": "versions",
"name": "[variables('imageRelease')]",
"location": "[variables('location')]",
"dependsOn": [
  "[variables('imageNameGen2')]"
],
"properties": {
  "publishingProfile": {
    "storageAccountType": "Standard_LRS",
    "targetRegions": [
      {
        "name": "[variables('location')]",
        "regionalReplicaCount": "1"
      }
    ]
  },
  "storageProfile": {
    "osDiskImage": {
      "source": {
        "id": "[resourceId('Microsoft.Storage/storageAccounts',
parameters('storageAccount'))]",
        "uri": "[parameters('vhdBlobURL')]"
      }
    }
  }
}
]
```

7.10.11. Networking requirements for user-provisioned infrastructure

All the Red Hat Enterprise Linux CoreOS (RHCOS) machines require networking to be configured in **initramfs** during boot to fetch their Ignition config files.

7.10.11.1. Network connectivity requirements

You must configure the network connectivity between machines to allow OpenShift Container Platform cluster components to communicate. Each machine must be able to resolve the hostnames of all other machines in the cluster.

This section provides details about the ports that are required.



IMPORTANT

In connected OpenShift Container Platform environments, all nodes are required to have internet access to pull images for platform containers and provide telemetry data to Red Hat.

Table 7.22. Ports used for all-machine to all-machine communications

Protocol	Port	Description
ICMP	N/A	Network reachability tests
TCP	1936	Metrics
	9000-9999	Host level services, including the node exporter on ports 9100-9101 and the Cluster Version Operator on port 9099 .
	10250-10259	The default ports that Kubernetes reserves
UDP	4789	VXLAN
	6081	Geneve
	9000-9999	Host level services, including the node exporter on ports 9100-9101 .
	500	IPsec IKE packets
	4500	IPsec NAT-T packets
	123	Network Time Protocol (NTP) on UDP port 123 If an external NTP time server is configured, you must open UDP port 123 .
TCP/UDP	30000-32767	Kubernetes node port
ESP	N/A	IPsec Encapsulating Security Payload (ESP)

Table 7.23. Ports used for all-machine to control plane communications

Protocol	Port	Description
TCP	6443	Kubernetes API

Table 7.24. Ports used for control plane machine to control plane machine communications

Protocol	Port	Description
TCP	2379-2380	etcd server and peer ports

7.10.12. Creating networking and load balancing components in Azure

You must configure networking and load balancing in Microsoft Azure for your OpenShift Container Platform cluster to use. One way to create these components is to modify the provided Azure Resource Manager (ARM) template.



NOTE

If you do not use the provided ARM template to create your Azure infrastructure, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- Configure an Azure account.
- Generate the Ignition config files for your cluster.
- Create and configure a VNet and associated subnets in Azure.

Procedure

1. Copy the template from the **ARM template for the network and load balancers** section of this topic and save it as **03_infra.json** in your cluster's installation directory. This template describes the networking and load balancing objects that your cluster requires.
2. Create the deployment by using the **az** CLI:

```
$ az deployment group create -g ${RESOURCE_GROUP} \
  --template-file "<installation_directory>/03_infra.json" \
  --parameters privateDNSZoneName="${CLUSTER_NAME}.${BASE_DOMAIN}" \ 1
  --parameters baseName="${INFRA_ID}" 2
```

1 The name of the private DNS zone.

2 The base name to be used in resource names; this is usually the cluster's infrastructure ID.

3. Create an **api** DNS record in the public zone for the API public load balancer. The **\${BASE_DOMAIN_RESOURCE_GROUP}** variable must point to the resource group where the public DNS zone exists.
 - a. Export the following variable:

```
$ export PUBLIC_IP=`az network public-ip list -g ${RESOURCE_GROUP} --query "[?
  name=='${INFRA_ID}-master-pip'] | [0].ipAddress" -o tsv`
```

- b. Create the **api** DNS record in a new public zone:

```
$ az network dns record-set a add-record -g ${BASE_DOMAIN_RESOURCE_GROUP} -
  z ${CLUSTER_NAME}.${BASE_DOMAIN} -n api -a ${PUBLIC_IP} --ttl 60
```

If you are adding the cluster to an existing public zone, you can create the **api** DNS record in it instead:

```
$ az network dns record-set a add-record -g ${BASE_DOMAIN_RESOURCE_GROUP} -
z ${BASE_DOMAIN} -n api.${CLUSTER_NAME} -a ${PUBLIC_IP} --ttl 60
```

7.10.12.1. ARM template for the network and load balancers

You can use the following Azure Resource Manager (ARM) template to deploy the networking objects and load balancers that you need for your OpenShift Container Platform cluster:

Example 7.60. 03_infra.json ARM template

```
{
  "$schema" : "https://schema.management.azure.com/schemas/2015-01-01/deploymentTemplate.json#",
  "contentVersion" : "1.0.0.0",
  "parameters" : {
    "baseName" : {
      "type" : "string",
      "minLength" : 1,
      "metadata" : {
        "description" : "Base name to be used in resource names (usually the cluster's Infra ID)"
      }
    },
    "vnetBaseName": {
      "type": "string",
      "defaultValue": "",
      "metadata" : {
        "description" : "The specific customer vnet's base name (optional)"
      }
    },
    "privateDNSZoneName" : {
      "type" : "string",
      "metadata" : {
        "description" : "Name of the private DNS zone"
      }
    }
  },
  "variables" : {
    "location" : "[resourceGroup().location]",
    "virtualNetworkName" : "[concat(if(not(empty(parameters('vnetBaseName'))),
parameters('vnetBaseName'), parameters('baseName')), '-vnet')]",
    "virtualNetworkID" : "[resourceId('Microsoft.Network/virtualNetworks',
variables('virtualNetworkName'))]",
    "masterSubnetName" : "[concat(if(not(empty(parameters('vnetBaseName'))),
parameters('vnetBaseName'), parameters('baseName')), '-master-subnet')]",
    "masterSubnetRef" : "[concat(variables('virtualNetworkID'), '/subnets/',
variables('masterSubnetName'))]",
    "masterPublicIpAddressName" : "[concat(parameters('baseName'), '-master-pip')]",
    "masterPublicIpAddressID" : "[resourceId('Microsoft.Network/publicIPAddresses',
variables('masterPublicIpAddressName'))]",
    "masterLoadBalancerName" : "[parameters('baseName')]",
    "masterLoadBalancerID" : "[resourceId('Microsoft.Network/loadBalancers',
variables('masterLoadBalancerName'))]",
    "internalLoadBalancerName" : "[concat(parameters('baseName'), '-internal-lb')]",
    "internalLoadBalancerID" : "[resourceId('Microsoft.Network/loadBalancers',
variables('internalLoadBalancerName'))]",
```

```

"skuName": "Standard"
},
"resources" : [
{
  "apiVersion" : "2018-12-01",
  "type" : "Microsoft.Network/publicIPAddresses",
  "name" : "[variables('masterPublicIpAddressName')]",
  "location" : "[variables('location')]",
  "sku": {
    "name": "[variables('skuName')]"
  },
  "properties" : {
    "publicIPAllocationMethod" : "Static",
    "dnsSettings" : {
      "domainNameLabel" : "[variables('masterPublicIpAddressName')]"
    }
  }
},
{
  "apiVersion" : "2018-12-01",
  "type" : "Microsoft.Network/loadBalancers",
  "name" : "[variables('masterLoadBalancerName')]",
  "location" : "[variables('location')]",
  "sku": {
    "name": "[variables('skuName')]"
  },
  "dependsOn" : [
    "[concat('Microsoft.Network/publicIPAddresses/', variables('masterPublicIpAddressName'))]"
  ],
  "properties" : {
    "frontendIPConfigurations" : [
      {
        "name" : "public-lb-ip-v4",
        "properties" : {
          "publicIPAddress" : {
            "id" : "[variables('masterPublicIpAddressID')]"
          }
        }
      }
    ],
    "backendAddressPools" : [
      {
        "name" : "[variables('masterLoadBalancerName')]"
      }
    ],
    "loadBalancingRules" : [
      {
        "name" : "api-internal",
        "properties" : {
          "frontendIPConfiguration" : {
            "id" : "[concat(variables('masterLoadBalancerID'), '/frontendIPConfigurations/public-lb-ip-
v4')]"
          },
          "backendAddressPool" : {
            "id" : "[concat(variables('masterLoadBalancerID'), '/backendAddressPools/',
variables('masterLoadBalancerName'))]"
          }
        }
      }
    ]
  }
}
]
}

```



```

    },
    "protocol" : "Tcp",
    "loadDistribution" : "Default",
    "idleTimeoutInMinutes" : 30,
    "frontendPort" : 6443,
    "backendPort" : 6443,
    "probe" : {
      "id" : "[concat(variables('masterLoadBalancerID'), '/probes/api-internal-probe')]"
    }
  }
],
"probes" : [
  {
    "name" : "api-internal-probe",
    "properties" : {
      "protocol" : "Https",
      "port" : 6443,
      "requestPath" : "/readyz",
      "intervalInSeconds" : 10,
      "numberOfProbes" : 3
    }
  }
]
},
{
  "apiVersion" : "2018-12-01",
  "type" : "Microsoft.Network/loadBalancers",
  "name" : "[variables('internalLoadBalancerName')]",
  "location" : "[variables('location')]",
  "sku" : {
    "name" : "[variables('skuName')]"
  },
  "properties" : {
    "frontendIPConfigurations" : [
      {
        "name" : "internal-lb-ip",
        "properties" : {
          "privateIPAllocationMethod" : "Dynamic",
          "subnet" : {
            "id" : "[variables('masterSubnetRef')]"
          },
          "privateIPAddressVersion" : "IPv4"
        }
      }
    ],
    "backendAddressPools" : [
      {
        "name" : "internal-lb-backend"
      }
    ],
    "loadBalancingRules" : [
      {
        "name" : "api-internal",
        "properties" : {

```

```

    "frontendIPConfiguration" : {
      "id" : "[concat(variables('internalLoadBalancerID'), '/frontendIPConfigurations/internal-lb-
ip')]"
    },
    "frontendPort" : 6443,
    "backendPort" : 6443,
    "enableFloatingIP" : false,
    "idleTimeoutInMinutes" : 30,
    "protocol" : "Tcp",
    "enableTcpReset" : false,
    "loadDistribution" : "Default",
    "backendAddressPool" : {
      "id" : "[concat(variables('internalLoadBalancerID'), '/backendAddressPools/internal-lb-
backend')]"
    },
    "probe" : {
      "id" : "[concat(variables('internalLoadBalancerID'), '/probes/api-internal-probe')]"
    }
  },
  {
    "name" : "sint",
    "properties" : {
      "frontendIPConfiguration" : {
        "id" : "[concat(variables('internalLoadBalancerID'), '/frontendIPConfigurations/internal-lb-
ip')]"
      },
      "frontendPort" : 22623,
      "backendPort" : 22623,
      "enableFloatingIP" : false,
      "idleTimeoutInMinutes" : 30,
      "protocol" : "Tcp",
      "enableTcpReset" : false,
      "loadDistribution" : "Default",
      "backendAddressPool" : {
        "id" : "[concat(variables('internalLoadBalancerID'), '/backendAddressPools/internal-lb-
backend')]"
      },
      "probe" : {
        "id" : "[concat(variables('internalLoadBalancerID'), '/probes/sint-probe')]"
      }
    }
  }
],
"probes" : [
  {
    "name" : "api-internal-probe",
    "properties" : {
      "protocol" : "Https",
      "port" : 6443,
      "requestPath" : "/readyz",
      "intervalInSeconds" : 10,
      "numberOfProbes" : 3
    }
  }
],
{

```


```

    "name" : "sint-probe",
    "properties" : {
      "protocol" : "Https",
      "port" : 22623,
      "requestPath" : "/healthz",
      "intervalInSeconds" : 10,
      "numberOfProbes" : 3
    }
  }
]
},
{
  "apiVersion": "2018-09-01",
  "type": "Microsoft.Network/privateDnsZones/A",
  "name": "[concat(parameters('privateDNSZoneName'), '/api')]",
  "location" : "[variables('location')]",
  "dependsOn" : [
    "[concat('Microsoft.Network/loadBalancers/', variables('internalLoadBalancerName'))]"
  ],
  "properties": {
    "ttl": 60,
    "aRecords": [
      {
        "ipv4Address": "[reference(variables('internalLoadBalancerName')).frontendIPConfigurations[0].properties.privateIPAddress]"
      }
    ]
  }
},
{
  "apiVersion": "2018-09-01",
  "type": "Microsoft.Network/privateDnsZones/A",
  "name": "[concat(parameters('privateDNSZoneName'), '/api-int')]",
  "location" : "[variables('location')]",
  "dependsOn" : [
    "[concat('Microsoft.Network/loadBalancers/', variables('internalLoadBalancerName'))]"
  ],
  "properties": {
    "ttl": 60,
    "aRecords": [
      {
        "ipv4Address": "[reference(variables('internalLoadBalancerName')).frontendIPConfigurations[0].properties.privateIPAddress]"
      }
    ]
  }
}
]
}

```

7.10.13. Creating the bootstrap machine in Azure

You must create the bootstrap machine in Microsoft Azure to use during OpenShift Container Platform cluster initialization. One way to create this machine is to modify the provided Azure Resource Manager (ARM) template.



NOTE

If you do not use the provided ARM template to create your bootstrap machine, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- Configure an Azure account.
- Generate the Ignition config files for your cluster.
- Create and configure a VNet and associated subnets in Azure.
- Create and configure networking and load balancers in Azure.
- Create control plane and compute roles.

Procedure

1. Copy the template from the **ARM template for the bootstrap machine** section of this topic and save it as **04_bootstrap.json** in your cluster's installation directory. This template describes the bootstrap machine that your cluster requires.

2. Export the bootstrap URL variable:

```
$ bootstrap_url_expiry=`date -u -d "10 hours" '+%Y-%m-%dT%H:%MZ`
```

```
$ export BOOTSTRAP_URL=`az storage blob generate-sas -c 'files' -n 'bootstrap.ign' --https-only --full-uri --permissions r --expiry $bootstrap_url_expiry --account-name ${CLUSTER_NAME}sa --account-key ${ACCOUNT_KEY} -o tsv`
```

3. Export the bootstrap ignition variable:

```
$ export BOOTSTRAP_IGNITION=`jq -rcnM --arg v "3.2.0" --arg url ${BOOTSTRAP_URL} '{ignition:{version:$v,config:{replace:{source:$url}}}}' | base64 | tr -d '\n`
```

4. Create the deployment by using the **az** CLI:

```
$ az deployment group create -g ${RESOURCE_GROUP} \
  --template-file "<installation_directory>/04_bootstrap.json" \
  --parameters bootstrapIgnition="${BOOTSTRAP_IGNITION}" \ 1
  --parameters baseName="${INFRA_ID}" \ 2
  --parameter bootstrapVMSize="Standard_D4s_v3" 3
```

1 The bootstrap Ignition content for the bootstrap cluster.

2 The base name to be used in resource names; this is usually the cluster's infrastructure ID.

- 3 Optional: Specify the size of the bootstrap VM. Use a VM size compatible with your specified architecture. If this value is not defined, the default value from the template is

7.10.13.1. ARM template for the bootstrap machine

You can use the following Azure Resource Manager (ARM) template to deploy the bootstrap machine that you need for your OpenShift Container Platform cluster:

Example 7.61. 04_bootstrap.json ARM template

```
{
  "$schema" : "https://schema.management.azure.com/schemas/2015-01-01/deploymentTemplate.json#",
  "contentVersion" : "1.0.0.0",
  "parameters" : {
    "baseName" : {
      "type" : "string",
      "minLength" : 1,
      "metadata" : {
        "description" : "Base name to be used in resource names (usually the cluster's Infra ID)"
      }
    },
    "vnetBaseName": {
      "type": "string",
      "defaultValue": "",
      "metadata" : {
        "description" : "The specific customer vnet's base name (optional)"
      }
    },
    "bootstrapIgnition" : {
      "type" : "string",
      "minLength" : 1,
      "metadata" : {
        "description" : "Bootstrap ignition content for the bootstrap cluster"
      }
    },
    "sshKeyData" : {
      "type" : "securestring",
      "defaultValue" : "Unused",
      "metadata" : {
        "description" : "Unused"
      }
    },
    "bootstrapVMSize" : {
      "type" : "string",
      "defaultValue" : "Standard_D4s_v3",
      "metadata" : {
        "description" : "The size of the Bootstrap Virtual Machine"
      }
    },
    "hyperVGen": {
      "type": "string",
      "metadata": {
        "description": "VM generation image to use"
      }
    }
  }
}
```

```

    },
    "defaultValue": "V2",
    "allowedValues": [
      "V1",
      "V2"
    ]
  }
},
"variables" : {
  "location" : "[resourceGroup().location]",
  "virtualNetworkName" : "[concat(if(not(empty(parameters('vnetBaseName'))),
parameters('vnetBaseName'), parameters('baseName')), '-vnet')]",
  "virtualNetworkID" : "[resourceId('Microsoft.Network/virtualNetworks',
variables('virtualNetworkName'))]",
  "masterSubnetName" : "[concat(if(not(empty(parameters('vnetBaseName'))),
parameters('vnetBaseName'), parameters('baseName')), '-master-subnet')]",
  "masterSubnetRef" : "[concat(variables('virtualNetworkID'), '/subnets/',
variables('masterSubnetName'))]",
  "masterLoadBalancerName" : "[parameters('baseName')]",
  "internalLoadBalancerName" : "[concat(parameters('baseName'), '-internal-lb')]",
  "sshKeyPath" : "/home/core/.ssh/authorized_keys",
  "identityName" : "[concat(parameters('baseName'), '-identity')]",
  "vmName" : "[concat(parameters('baseName'), '-bootstrap')]",
  "nicName" : "[concat(variables('vmName'), '-nic')]",
  "galleryName": "[concat('gallery_', replace(parameters('baseName'), '-', '_'))]",
  "imageName" : "[concat(parameters('baseName'), if(equals(parameters('hyperVGen'), 'V2'), '-
gen2', ''))]",
  "clusterNsgName" : "[concat(if(not(empty(parameters('vnetBaseName'))),
parameters('vnetBaseName'), parameters('baseName')), '-nsg')]",
  "sshPublicIpAddressName" : "[concat(variables('vmName'), '-ssh-pip')]"
},
"resources" : [
  {
    "apiVersion" : "2018-12-01",
    "type" : "Microsoft.Network/publicIPAddresses",
    "name" : "[variables('sshPublicIpAddressName')]",
    "location" : "[variables('location')]",
    "sku": {
      "name": "Standard"
    },
  },
  "properties" : {
    "publicIPAllocationMethod" : "Static",
    "dnsSettings" : {
      "domainNameLabel" : "[variables('sshPublicIpAddressName')]"
    }
  }
},
{
  "apiVersion" : "2018-06-01",
  "type" : "Microsoft.Network/networkInterfaces",
  "name" : "[variables('nicName')]",
  "location" : "[variables('location')]",
  "dependsOn" : [
    "[resourceId('Microsoft.Network/publicIPAddresses', variables('sshPublicIpAddressName'))]"
  ],
  "properties" : {

```

```

"ipConfigurations" : [
  {
    "name" : "pipConfig",
    "properties" : {
      "privateIPAllocationMethod" : "Dynamic",
      "publicIPAddress": {
        "id": "[resourceId('Microsoft.Network/publicIPAddresses',
variables('sshPublicIpAddressName'))]"
      },
      "subnet" : {
        "id" : "[variables('masterSubnetRef')]"
      },
      "loadBalancerBackendAddressPools" : [
        {
          "id" : "[concat('/subscriptions/', subscription().subscriptionId, '/resourceGroups/',
resourceGroup().name, '/providers/Microsoft.Network/loadBalancers/',
variables('masterLoadBalancerName'), '/backendAddressPools/',
variables('masterLoadBalancerName'))]"
        },
        {
          "id" : "[concat('/subscriptions/', subscription().subscriptionId, '/resourceGroups/',
resourceGroup().name, '/providers/Microsoft.Network/loadBalancers/',
variables('internalLoadBalancerName'), '/backendAddressPools/internal-lb-backend')]"
        }
      ]
    }
  }
]
},
{
  "apiVersion" : "2018-06-01",
  "type" : "Microsoft.Compute/virtualMachines",
  "name" : "[variables('vmName')]",
  "location" : "[variables('location')]",
  "identity" : {
    "type" : "userAssigned",
    "userAssignedIdentities" : {
      "[resourceID('Microsoft.ManagedIdentity/userAssignedIdentities/',
variables('identityName'))]" : {}
    }
  },
  "dependsOn" : [
    "[concat('Microsoft.Network/networkInterfaces/', variables('nicName'))]"
  ],
  "properties" : {
    "hardwareProfile" : {
      "vmSize" : "[parameters('bootstrapVMSize')]"
    },
    "osProfile" : {
      "computerName" : "[variables('vmName')]",
      "adminUsername" : "core",
      "adminPassword" : "NotActuallyApplied!",
      "customData" : "[parameters('bootstrapIgnition')]",
      "linuxConfiguration" : {
        "disablePasswordAuthentication" : false
      }
    }
  }
}
]
}
}

```

```

    }
  },
  "storageProfile" : {
    "imageReference": {
      "id": "[resourceId('Microsoft.Compute/galleries/images', variables('galleryName'),
variables('imageName'))]"
    },
    "osDisk" : {
      "name": "[concat(variables('vmName'),'_OSDisk')]",
      "osType" : "Linux",
      "createOption" : "FromImage",
      "managedDisk": {
        "storageAccountType": "Premium_LRS"
      },
      "diskSizeGB" : 100
    }
  },
  "networkProfile" : {
    "networkInterfaces" : [
      {
        "id": "[resourceId('Microsoft.Network/networkInterfaces', variables('nicName'))]"
      }
    ]
  }
},
{
  "apiVersion" : "2018-06-01",
  "type": "Microsoft.Network/networkSecurityGroups/securityRules",
  "name" : "[concat(variables('clusterNsgName'), '/bootstrap_ssh_in')]",
  "location" : "[variables('location')]",
  "dependsOn" : [
    "[resourceId('Microsoft.Compute/virtualMachines', variables('vmName'))]"
  ],
  "properties": {
    "protocol" : "Tcp",
    "sourcePortRange" : "*",
    "destinationPortRange" : "22",
    "sourceAddressPrefix" : "*",
    "destinationAddressPrefix" : "*",
    "access" : "Allow",
    "priority" : 100,
    "direction" : "Inbound"
  }
}
]
}

```

7.10.14. Creating the control plane machines in Azure

You must create the control plane machines in Microsoft Azure for your cluster to use. One way to create these machines is to modify the provided Azure Resource Manager (ARM) template.



NOTE

By default, Microsoft Azure places control plane machines and compute machines in a pre-set availability zone. You can manually set an availability zone for a compute node or control plane node. To do this, modify a vendor's Azure Resource Manager (ARM) template by specifying each of your availability zones in the **zones** parameter of the virtual machine resource.

If you do not use the provided ARM template to create your control plane machines, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, consider contacting Red Hat support with your installation logs.

Prerequisites

- Configure an Azure account.
- Generate the Ignition config files for your cluster.
- Create and configure a VNet and associated subnets in Azure.
- Create and configure networking and load balancers in Azure.
- Create control plane and compute roles.
- Create the bootstrap machine.

Procedure

1. Copy the template from the **ARM template for control plane machines** section of this topic and save it as **05_masters.json** in your cluster's installation directory. This template describes the control plane machines that your cluster requires.
2. Export the following variable needed by the control plane machine deployment:

```
$ export MASTER_IGNITION=`cat <installation_directory>/master.ign | base64 | tr -d '\n'`
```

3. Create the deployment by using the **az** CLI:

```
$ az deployment group create -g ${RESOURCE_GROUP} \
  --template-file "<installation_directory>/05_masters.json" \
  --parameters masterIgnition="${MASTER_IGNITION}" \ 1
  --parameters baseName="${INFRA_ID}" \ 2
  --parameters masterVMSize="Standard_D8s_v3" \ 3
```

- 1** The Ignition content for the control plane nodes.
- 2** The base name to be used in resource names; this is usually the cluster's infrastructure ID.
- 3** Optional: Specify the size of the Control Plane VM. Use a VM size compatible with your specified architecture. If this value is not defined, the default value from the template is set.

7.10.14.1. ARM template for control plane machines

You can use the following Azure Resource Manager (ARM) template to deploy the control plane machines that you need for your OpenShift Container Platform cluster:

Example 7.62. 05_masters.json ARM template

```
{
  "$schema" : "https://schema.management.azure.com/schemas/2015-01-01/deploymentTemplate.json#",
  "contentVersion" : "1.0.0.0",
  "parameters" : {
    "baseName" : {
      "type" : "string",
      "minLength" : 1,
      "metadata" : {
        "description" : "Base name to be used in resource names (usually the cluster's Infra ID)"
      }
    },
    "vnetBaseName" : {
      "type" : "string",
      "defaultValue" : "",
      "metadata" : {
        "description" : "The specific customer vnet's base name (optional)"
      }
    },
    "masterIgnition" : {
      "type" : "string",
      "metadata" : {
        "description" : "Ignition content for the master nodes"
      }
    },
    "numberOfMasters" : {
      "type" : "int",
      "defaultValue" : 3,
      "minValue" : 2,
      "maxValue" : 30,
      "metadata" : {
        "description" : "Number of OpenShift masters to deploy"
      }
    },
    "sshKeyData" : {
      "type" : "securestring",
      "defaultValue" : "Unused",
      "metadata" : {
        "description" : "Unused"
      }
    },
    "privateDNSZoneName" : {
      "type" : "string",
      "defaultValue" : "",
      "metadata" : {
        "description" : "unused"
      }
    },
    "masterVMSize" : {
      "type" : "string",
      "defaultValue" : "Standard_D8s_v3",
```

```

    "metadata" : {
      "description" : "The size of the Master Virtual Machines"
    }
  },
  "diskSizeGB" : {
    "type" : "int",
    "defaultValue" : 1024,
    "metadata" : {
      "description" : "Size of the Master VM OS disk, in GB"
    }
  },
  "hyperVGen": {
    "type": "string",
    "metadata": {
      "description": "VM generation image to use"
    },
    "defaultValue": "V2",
    "allowedValues": [
      "V1",
      "V2"
    ]
  },
  "variables" : {
    "location" : "[resourceGroup().location]",
    "virtualNetworkName" : "[concat(if(not(empty(parameters('vnetBaseName'))),
parameters('vnetBaseName'), parameters('baseName')), '-vnet')]",
    "virtualNetworkID" : "[resourceId('Microsoft.Network/virtualNetworks',
variables('virtualNetworkName'))]",
    "masterSubnetName" : "[concat(if(not(empty(parameters('vnetBaseName'))),
parameters('vnetBaseName'), parameters('baseName')), '-master-subnet')]",
    "masterSubnetRef" : "[concat(variables('virtualNetworkID'), '/subnets/',
variables('masterSubnetName'))]",
    "masterLoadBalancerName" : "[parameters('baseName')]",
    "internalLoadBalancerName" : "[concat(parameters('baseName'), '-internal-lb')]",
    "sshKeyPath" : "/home/core/.ssh/authorized_keys",
    "identityName" : "[concat(parameters('baseName'), '-identity')]",
    "galleryName": "[concat('gallery_', replace(parameters('baseName'), '-', '_'))]",
    "imageName" : "[concat(parameters('baseName'), if(equals(parameters('hyperVGen'), 'V2'), '-
gen2', ''))]",
    "copy" : [
      {
        "name" : "vmNames",
        "count" : "[parameters('numberOfMasters')]",
        "input" : "[concat(parameters('baseName'), '-master-', copyIndex('vmNames'))]"
      }
    ]
  },
  "resources" : [
    {
      "apiVersion" : "2018-06-01",
      "type" : "Microsoft.Network/networkInterfaces",
      "copy" : {
        "name" : "nicCopy",
        "count" : "[length(variables('vmNames'))]"
      }
    }
  ]
}

```

```

"name" : "[concat(variables('vmNames')[copyIndex()], '-nic')]",
"location" : "[variables('location')]",
"properties" : {
  "ipConfigurations" : [
    {
      "name" : "pipConfig",
      "properties" : {
        "privateIPAllocationMethod" : "Dynamic",
        "subnet" : {
          "id" : "[variables('masterSubnetRef')]"
        },
        "loadBalancerBackendAddressPools" : [
          {
            "id" : "[concat('/subscriptions/', subscription().subscriptionId, '/resourceGroups/',
resourceGroup().name, '/providers/Microsoft.Network/loadBalancers/',
variables('masterLoadBalancerName'), '/backendAddressPools/',
variables('masterLoadBalancerName'))]"
          },
          {
            "id" : "[concat('/subscriptions/', subscription().subscriptionId, '/resourceGroups/',
resourceGroup().name, '/providers/Microsoft.Network/loadBalancers/',
variables('internalLoadBalancerName'), '/backendAddressPools/internal-lb-backend')]"
          }
        ]
      }
    }
  ]
},
{
  "apiVersion" : "2018-06-01",
  "type" : "Microsoft.Compute/virtualMachines",
  "copy" : {
    "name" : "vmCopy",
    "count" : "[length(variables('vmNames'))]"
  },
  "name" : "[variables('vmNames')[copyIndex()]]",
  "location" : "[variables('location')]",
  "identity" : {
    "type" : "userAssigned",
    "userAssignedIdentities" : {
      "[resourceID('Microsoft.ManagedIdentity/userAssignedIdentities/',
variables('identityName'))]" : {}
    }
  },
  "dependsOn" : [
    "[concat('Microsoft.Network/networkInterfaces/', concat(variables('vmNames')[copyIndex()], '-
nic'))]"
  ],
  "properties" : {
    "hardwareProfile" : {
      "vmSize" : "[parameters('masterVMSize')]"
    },
    "osProfile" : {
      "computerName" : "[variables('vmNames')[copyIndex()]]",
      "adminUsername" : "core",

```

```

    "adminPassword" : "NotActuallyApplied!",
    "customData" : "[parameters('masterIgnition')]",
    "linuxConfiguration" : {
      "disablePasswordAuthentication" : false
    }
  },
  "storageProfile" : {
    "imageReference": {
      "id": "[resourceId('Microsoft.Compute/galleries/images', variables('galleryName'),
variables('imageName'))]"
    },
    "osDisk" : {
      "name": "[concat(variables('vmNames')[copyIndex()], '_OSDisk')]",
      "osType" : "Linux",
      "createOption" : "FromImage",
      "caching": "ReadOnly",
      "writeAcceleratorEnabled": false,
      "managedDisk": {
        "storageAccountType": "Premium_LRS"
      },
      "diskSizeGB" : "[parameters('diskSizeGB')]"
    }
  },
  "networkProfile" : {
    "networkInterfaces" : [
      {
        "id" : "[resourceId('Microsoft.Network/networkInterfaces', concat(variables('vmNames')
[copyIndex()], '-nic'))]",
        "properties": {
          "primary": false
        }
      }
    ]
  }
}
]
}

```

7.10.15. Wait for bootstrap completion and remove bootstrap resources in Azure

After you create all of the required infrastructure in Microsoft Azure, wait for the bootstrap process to complete on the machines that you provisioned by using the Ignition config files that you generated with the installation program.

Prerequisites

- Configure an Azure account.
- Generate the Ignition config files for your cluster.
- Create and configure a VNet and associated subnets in Azure.
- Create and configure networking and load balancers in Azure.

- Create control plane and compute roles.
- Create the bootstrap machine.
- Create the control plane machines.

Procedure

1. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install wait-for bootstrap-complete --dir <installation_directory> \ 1
--log-level info 2
```

1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

If the command exits without a **FATAL** warning, your production control plane has initialized.

2. Delete the bootstrap resources:

```
$ az network nsg rule delete -g ${RESOURCE_GROUP} --nsg-name ${INFRA_ID}-nsg --
name bootstrap_ssh_in
$ az vm stop -g ${RESOURCE_GROUP} --name ${INFRA_ID}-bootstrap
$ az vm deallocate -g ${RESOURCE_GROUP} --name ${INFRA_ID}-bootstrap
$ az vm delete -g ${RESOURCE_GROUP} --name ${INFRA_ID}-bootstrap --yes
$ az disk delete -g ${RESOURCE_GROUP} --name ${INFRA_ID}-bootstrap OSDisk --no-
wait --yes
$ az network nic delete -g ${RESOURCE_GROUP} --name ${INFRA_ID}-bootstrap-nic --no-
wait
$ az storage blob delete --account-key ${ACCOUNT_KEY} --account-name
${CLUSTER_NAME}sa --container-name files --name bootstrap.ign
$ az network public-ip delete -g ${RESOURCE_GROUP} --name ${INFRA_ID}-bootstrap-
ssh-pip
```



NOTE

If you do not delete the bootstrap server, installation may not succeed due to API traffic being routed to the bootstrap server.

7.10.16. Creating additional worker machines in Azure

You can create worker machines in Microsoft Azure for your cluster to use by launching individual instances discretely or by automated processes outside the cluster, such as auto scaling groups. You can also take advantage of the built-in cluster scaling mechanisms and the machine API in OpenShift Container Platform.

In this example, you manually launch one instance by using the Azure Resource Manager (ARM) template. Additional instances can be launched by including additional resources of type **06_workers.json** in the file.



NOTE

By default, Microsoft Azure places control plane machines and compute machines in a pre-set availability zone. You can manually set an availability zone for a compute node or control plane node. To do this, modify a vendor's ARM template by specifying each of your availability zones in the **zones** parameter of the virtual machine resource.

If you do not use the provided ARM template to create your control plane machines, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, consider contacting Red Hat support with your installation logs.

Prerequisites

- Configure an Azure account.
- Generate the Ignition config files for your cluster.
- Create and configure a VNet and associated subnets in Azure.
- Create and configure networking and load balancers in Azure.
- Create control plane and compute roles.
- Create the bootstrap machine.
- Create the control plane machines.

Procedure

1. Copy the template from the **ARM template for worker machines** section of this topic and save it as **06_workers.json** in your cluster's installation directory. This template describes the worker machines that your cluster requires.
2. Export the following variable needed by the worker machine deployment:

```
$ export WORKER_IGNITION=`cat <installation_directory>/worker.ign | base64 | tr -d '\n'`
```

3. Create the deployment by using the **az** CLI:

```
$ az deployment group create -g ${RESOURCE_GROUP} \
  --template-file "<installation_directory>/06_workers.json" \
  --parameters workerIgnition="${WORKER_IGNITION}" \ 1
  --parameters baseName="${INFRA_ID}" \ 2
  --parameters nodeVMSize="Standard_D4s_v3" \ 3
```

- 1** The Ignition content for the worker nodes.
- 2** The base name to be used in resource names; this is usually the cluster's infrastructure ID.
- 3** Optional: Specify the size of the compute node VM. Use a VM size compatible with your specified architecture. If this value is not defined, the default value from the template is set.

7.10.16.1. ARM template for worker machines

You can use the following Azure Resource Manager (ARM) template to deploy the worker machines that you need for your OpenShift Container Platform cluster:

Example 7.63. 06_workers.json ARM template

```
{
  "$schema" : "https://schema.management.azure.com/schemas/2015-01-01/deploymentTemplate.json#",
  "contentVersion" : "1.0.0.0",
  "parameters" : {
    "baseName" : {
      "type" : "string",
      "minLength" : 1,
      "metadata" : {
        "description" : "Base name to be used in resource names (usually the cluster's Infra ID)"
      }
    },
    "vnetBaseName": {
      "type": "string",
      "defaultValue": "",
      "metadata" : {
        "description" : "The specific customer vnet's base name (optional)"
      }
    },
    "workerIgnition" : {
      "type" : "string",
      "metadata" : {
        "description" : "Ignition content for the worker nodes"
      }
    },
    "numberOfNodes" : {
      "type" : "int",
      "defaultValue" : 3,
      "minValue" : 2,
      "maxValue" : 30,
      "metadata" : {
        "description" : "Number of OpenShift compute nodes to deploy"
      }
    },
    "sshKeyData" : {
      "type" : "securestring",
      "defaultValue" : "Unused",
      "metadata" : {
        "description" : "Unused"
      }
    },
    "nodeVMSize" : {
      "type" : "string",
      "defaultValue" : "Standard_D4s_v3",
      "metadata" : {
        "description" : "The size of the each Node Virtual Machine"
      }
    },
    "hyperVGen": {
```



```

    "type": "string",
    "metadata": {
      "description": "VM generation image to use"
    },
    "defaultValue": "V2",
    "allowedValues": [
      "V1",
      "V2"
    ]
  }
},
"variables": {
  "location": "[resourceGroup().location]",
  "virtualNetworkName" : "[concat(if(not(empty(parameters('vnetBaseName'))),
parameters('vnetBaseName'), parameters('baseName')), '-vnet')]",
  "virtualNetworkID" : "[resourceId('Microsoft.Network/virtualNetworks',
variables('virtualNetworkName'))]",
  "nodeSubnetName" : "[concat(if(not(empty(parameters('vnetBaseName'))),
parameters('vnetBaseName'), parameters('baseName')), '-worker-subnet')]",
  "nodeSubnetRef" : "[concat(variables('virtualNetworkID'), '/subnets/',
variables('nodeSubnetName'))]",
  "infraLoadBalancerName" : "[parameters('baseName')]",
  "sshKeyPath" : "/home/capi/.ssh/authorized_keys",
  "identityName" : "[concat(parameters('baseName'), '-identity')]",
  "galleryName": "[concat('gallery_', replace(parameters('baseName'), '-', '_'))]",
  "imageName" : "[concat(parameters('baseName'), if(equals(parameters('hyperVGen'), 'V2'), '-
gen2', ''))]",
  "copy" : [
    {
      "name" : "vmNames",
      "count" : "[parameters('numberOfNodes')]",
      "input" : "[concat(parameters('baseName'), '-worker-', variables('location'), '-',
copyIndex('vmNames', 1))]"
    }
  ]
},
"resources" : [
  {
    "apiVersion" : "2019-05-01",
    "name" : "[concat('node', copyIndex())]",
    "type" : "Microsoft.Resources/deployments",
    "copy" : {
      "name" : "nodeCopy",
      "count" : "[length(variables('vmNames'))]"
    },
    "properties" : {
      "mode" : "Incremental",
      "template" : {
        "$schema" : "http://schema.management.azure.com/schemas/2015-01-
01/deploymentTemplate.json#",
        "contentVersion" : "1.0.0.0",
        "resources" : [
          {
            "apiVersion" : "2018-06-01",
            "type" : "Microsoft.Network/networkInterfaces",
            "name" : "[concat(variables('vmNames')[copyIndex()], '-nic')]",

```

```

"location" : "[variables('location')]",
"properties" : {
  "ipConfigurations" : [
    {
      "name" : "pipConfig",
      "properties" : {
        "privateIPAllocationMethod" : "Dynamic",
        "subnet" : {
          "id" : "[variables('nodeSubnetRef')]"
        }
      }
    }
  ]
},
{
  "apiVersion" : "2018-06-01",
  "type" : "Microsoft.Compute/virtualMachines",
  "name" : "[variables('vmNames')[copyIndex()]]",
  "location" : "[variables('location')]",
  "tags" : {
    "kubernetes.io-cluster-ffranzupi": "owned"
  },
  "identity" : {
    "type" : "userAssigned",
    "userAssignedIdentities" : {
      "[resourceID('Microsoft.ManagedIdentity/userAssignedIdentities/',
variables('identityName'))]" : {}
    }
  },
  "dependsOn" : [
    "[concat('Microsoft.Network/networkInterfaces/', concat(variables('vmNames')
[copyIndex()], '-nic'))]"
  ],
  "properties" : {
    "hardwareProfile" : {
      "vmSize" : "[parameters('nodeVMSize')]"
    },
    "osProfile" : {
      "computerName" : "[variables('vmNames')[copyIndex()]]",
      "adminUsername" : "capi",
      "adminPassword" : "NotActuallyApplied!",
      "customData" : "[parameters('workerIgnition')]",
      "linuxConfiguration" : {
        "disablePasswordAuthentication" : false
      }
    },
    "storageProfile" : {
      "imageReference" : {
        "id" : "[resourceID('Microsoft.Compute/galleries/images', variables('galleryName'),
variables('imageName'))]"
      },
      "osDisk" : {
        "name" : "[concat(variables('vmNames')[copyIndex()], '_OSDisk')]",
        "osType" : "Linux",
        "createOption" : "FromImage",

```


```

    "managedDisk": {
      "storageAccountType": "Premium_LRS"
    },
    "diskSizeGB": 128
  },
  "networkProfile": {
    "networkInterfaces": [
      {
        "id": "[resourceId('Microsoft.Network/networkInterfaces',
concat(variables('vmNames')[copyIndex()], '-nic'))]",
        "properties": {
          "primary": true
        }
      }
    ]
  }
}
]
}
}
]
}
}

```

7.10.17. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.



IMPORTANT

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**.
To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.



NOTE

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.

5. Move the **oc** binary to a directory on your PATH.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

7.10.18. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

7.10.19. Approving the certificate signing requests for your machines

When you add machines to a cluster, two pending certificate signing requests (CSRs) are generated for each machine that you added. You must confirm that these CSRs are approved or, if necessary, approve them yourself. The client requests must be approved first, followed by the server requests.

Prerequisites

- You added machines to your cluster.

Procedure

1. Confirm that the cluster recognizes the machines:

```
$ oc get nodes
```

Example output

```
NAME      STATUS   ROLES    AGE   VERSION
master-0  Ready   master   63m   v1.29.4
master-1  Ready   master   63m   v1.29.4
master-2  Ready   master   64m   v1.29.4
```

The output lists all of the machines that you created.



NOTE

The preceding output might not include the compute nodes, also known as worker nodes, until some CSRs are approved.

2. Review the pending CSRs and ensure that you see the client requests with the **Pending** or **Approved** status for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

```
NAME      AGE   REQUESTOR                                     CONDITION
csr-8b2br  15m   system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper  Pending
csr-8vnps  15m   system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper  Pending
...
```

In this example, two machines are joining the cluster. You might see more approved CSRs in the list.

3. If the CSRs were not approved, after all of the pending CSRs for the machines you added are in **Pending** status, approve the CSRs for your cluster machines:



NOTE

Because the CSRs rotate automatically, approve your CSRs within an hour of adding the machines to the cluster. If you do not approve them within an hour, the certificates will rotate, and more than two certificates will be present for each node. You must approve all of these certificates. After the client CSR is approved, the Kubelet creates a secondary CSR for the serving certificate, which requires manual approval. Then, subsequent serving certificate renewal requests are automatically approved by the **machine-approver** if the Kubelet requests a new certificate with identical parameters.

**NOTE**

For clusters running on platforms that are not machine API enabled, such as bare metal and other user-provisioned infrastructure, you must implement a method of automatically approving the kubelet serving certificate requests (CSRs). If a request is not approved, then the **oc exec**, **oc rsh**, and **oc logs** commands cannot succeed, because a serving certificate is required when the API server connects to the kubelet. Any operation that contacts the Kubelet endpoint requires this certificate approval to be in place. The method must watch for new CSRs, confirm that the CSR was submitted by the **node-bootstrap** service account in the **system:node** or **system:admin** groups, and confirm the identity of the node.

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{"\n"}{{end}}{{end}}' | xargs --no-run-if-empty oc adm certificate approve
```

**NOTE**

Some Operators might not become available until some CSRs are approved.

4. Now that your client requests are approved, you must review the server requests for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

```
NAME      AGE   REQUESTOR                                     CONDITION
csr-bfd72 5m26s system:node:ip-10-0-50-126.us-east-2.compute.internal
Pending
csr-c57lv 5m26s system:node:ip-10-0-95-157.us-east-2.compute.internal
Pending
...
```

5. If the remaining CSRs are not approved, and are in the **Pending** status, approve the CSRs for your cluster machines:

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{\n"}\n{{end}}\n{{end}}' | xargs oc adm certificate approve
```

- After all client and server CSRs have been approved, the machines have the **Ready** status. Verify this by running the following command:

```
$ oc get nodes
```

Example output

```
NAME      STATUS  ROLES  AGE  VERSION
master-0  Ready   master 73m  v1.29.4
master-1  Ready   master 73m  v1.29.4
master-2  Ready   master 74m  v1.29.4
worker-0  Ready   worker 11m  v1.29.4
worker-1  Ready   worker 11m  v1.29.4
```



NOTE

It can take a few minutes after approval of the server CSRs for the machines to transition to the **Ready** status.

Additional information

- For more information on CSRs, see [Certificate Signing Requests](#).

7.10.20. Adding the Ingress DNS records

If you removed the DNS Zone configuration when creating Kubernetes manifests and generating Ignition configs, you must manually create DNS records that point at the Ingress load balancer. You can create either a wildcard ***.apps.{baseDomain}**, or specific records. You can use A, CNAME, and other records per your requirements.

Prerequisites

- You deployed an OpenShift Container Platform cluster on Microsoft Azure by using infrastructure that you provisioned.
- Install the OpenShift CLI (**oc**).
- Install or update the [Azure CLI](#).

Procedure

- Confirm the Ingress router has created a load balancer and populated the **EXTERNAL-IP** field:

```
$ oc -n openshift-ingress get service router-default
```

Example output


```

NAME          TYPE          CLUSTER-IP  EXTERNAL-IP  PORT(S)          AGE
router-default LoadBalancer 172.30.20.10 35.130.120.110
80:32288/TCP,443:31215/TCP 20

```

2. Export the Ingress router IP as a variable:

```
$ export PUBLIC_IP_ROUTER=`oc -n openshift-ingress get service router-default --no-headers | awk '{print $4}'`
```

3. Add a ***.apps** record to the public DNS zone.

- a. If you are adding this cluster to a new public zone, run:

```
$ az network dns record-set a add-record -g ${BASE_DOMAIN_RESOURCE_GROUP} -z ${CLUSTER_NAME}.${BASE_DOMAIN} -n *.apps -a ${PUBLIC_IP_ROUTER} --ttl 300
```

- b. If you are adding this cluster to an already existing public zone, run:

```
$ az network dns record-set a add-record -g ${BASE_DOMAIN_RESOURCE_GROUP} -z ${BASE_DOMAIN} -n *.apps.${CLUSTER_NAME} -a ${PUBLIC_IP_ROUTER} --ttl 300
```

4. Add a ***.apps** record to the private DNS zone:

- a. Create a ***.apps** record by using the following command:

```
$ az network private-dns record-set a create -g ${RESOURCE_GROUP} -z ${CLUSTER_NAME}.${BASE_DOMAIN} -n *.apps --ttl 300
```

- b. Add the ***.apps** record to the private DNS zone by using the following command:

```
$ az network private-dns record-set a add-record -g ${RESOURCE_GROUP} -z ${CLUSTER_NAME}.${BASE_DOMAIN} -n *.apps -a ${PUBLIC_IP_ROUTER}
```

If you prefer to add explicit domains instead of using a wildcard, you can create entries for each of the cluster's current routes:

```
$ oc get --all-namespaces -o jsonpath='{range .items[*]}{range .status.ingress[*]}{.host}{"\n"}{end}{end}' routes
```

Example output

```

oauth-openshift.apps.cluster.basedomain.com
console-openshift-console.apps.cluster.basedomain.com
downloads-openshift-console.apps.cluster.basedomain.com
alertmanager-main-openshift-monitoring.apps.cluster.basedomain.com
prometheus-k8s-openshift-monitoring.apps.cluster.basedomain.com

```

7.10.21. Completing an Azure installation on user-provisioned infrastructure

After you start the OpenShift Container Platform installation on Microsoft Azure user-provisioned infrastructure, you can monitor the cluster events until the cluster is ready.

Prerequisites

- Deploy the bootstrap machine for an OpenShift Container Platform cluster on user-provisioned Azure infrastructure.
- Install the **oc** CLI and log in.

Procedure

- Complete the cluster installation:

```
$ ./openshift-install --dir <installation_directory> wait-for install-complete 1
```

Example output

```
INFO Waiting up to 30m0s for the cluster to initialize...
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

7.10.22. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service

7.11. INSTALLING A CLUSTER ON AZURE USING ARM TEMPLATES

In OpenShift Container Platform version 4.16, you can install a cluster on Microsoft Azure by using infrastructure that you provide.

Several [Azure Resource Manager](#) (ARM) templates are provided to assist in completing these steps or to help model your own.




IMPORTANT

The steps for performing a user-provisioned infrastructure installation are provided as an example only. Installing a cluster with infrastructure you provide requires knowledge of the cloud provider and the installation process of OpenShift Container Platform. Several ARM templates are provided to assist in completing these steps or to help model your own. You are also free to create the required resources through other methods; the templates are just an example.

7.11.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [configured an Azure account](#) to host the cluster.
- You downloaded the Azure CLI and installed it on your computer. See [Install the Azure CLI](#) in the Azure documentation. The documentation below was last tested using version **2.38.0** of the Azure CLI. Azure CLI commands might perform differently based on the version you use.
- If the cloud identity and access management (IAM) APIs are not accessible in your environment, or if you do not want to store an administrator-level credential secret in the **kube-system** namespace, see [Alternatives to storing administrator-level secrets in the kube-system project](#).
- If you use a firewall and plan to use the Telemetry service, you [configured the firewall to allow the sites](#) that your cluster requires access to.



NOTE

Be sure to also review this site list if you are configuring a proxy.

7.11.2. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:


- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.

**IMPORTANT**

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

7.11.3. Configuring your Azure project


Before you can install OpenShift Container Platform, you must configure an Azure project to host it.

**IMPORTANT**

All Azure resources that are available through public endpoints are subject to resource name restrictions, and you cannot create resources that use certain terms. For a list of terms that Azure restricts, see [Resolve reserved resource name errors](#) in the Azure documentation.

7.11.3.1. Azure account limits

The OpenShift Container Platform cluster uses a number of Microsoft Azure components, and the default [Azure subscription and service limits, quotas, and constraints](#) affect your ability to install OpenShift Container Platform clusters.

**IMPORTANT**

Default limits vary by offer category types, such as Free Trial and Pay-As-You-Go, and by series, such as Dv2, F, and G. For example, the default for Enterprise Agreement subscriptions is 350 cores.


Check the limits for your subscription type and if necessary, increase quota limits for your account before you install a default cluster on Azure.

The following table summarizes the Azure components whose limits can impact your ability to install and run OpenShift Container Platform clusters.

Component	Number of components required by default	Default Azure limit	Description

Component	Number of components required by default	Default Azure limit	Description
vCPU	40	20 per region	<p>A default cluster requires 40 vCPUs, so you must increase the account limit.</p> <p>By default, each cluster creates the following instances:</p> <ul style="list-style-type: none"> • One bootstrap machine, which is removed after installation • Three control plane machines • Three compute machines <p>Because the bootstrap machine uses Standard_D4s_v3 machines, which use 4 vCPUs, the control plane machines use Standard_D8s_v3 virtual machines, which use 8 vCPUs, and the worker machines use Standard_D4s_v3 virtual machines, which use 4 vCPUs, a default cluster requires 40 vCPUs. The bootstrap node VM, which uses 4 vCPUs, is used only during installation.</p> <p>To deploy more worker nodes, enable autoscaling, deploy large workloads, or use a different instance type, you must further increase the vCPU limit for your account to ensure that your cluster can deploy the machines that you require.</p>
OS Disk	7		<p>Each cluster machine must have a minimum of 100 GB of storage and 300 IOPS. While these are the minimum supported values, faster storage is recommended for production clusters and clusters with intensive workloads. For more information about optimizing storage for performance, see the page titled "Optimizing storage" in the "Scalability and performance" section.</p>
VNet	1	1000 per region	<p>Each default cluster requires one Virtual Network (VNet), which contains two subnets.</p>
Network interfaces	7	65,536 per region	<p>Each default cluster requires seven network interfaces. If you create more machines or your deployed workloads create load balancers, your cluster uses more network interfaces.</p>

Component	Number of components required by default	Default Azure limit	Description						
Network security groups	2	5000	<p>Each cluster creates network security groups for each subnet in the VNet. The default cluster creates network security groups for the control plane and for the compute node subnets:</p> <table border="1"> <tr> <td>control plane</td> <td>Allows the control plane machines to be reached on port 6443 from anywhere</td> </tr> <tr> <td>node</td> <td>Allows worker nodes to be reached from the internet on ports 80 and 443</td> </tr> </table>	control plane	Allows the control plane machines to be reached on port 6443 from anywhere	node	Allows worker nodes to be reached from the internet on ports 80 and 443		
control plane	Allows the control plane machines to be reached on port 6443 from anywhere								
node	Allows worker nodes to be reached from the internet on ports 80 and 443								
Network load balancers	3	1000 per region	<p>Each cluster creates the following load balancers:</p> <table border="1"> <tr> <td>default</td> <td>Public IP address that load balances requests to ports 80 and 443 across worker machines</td> </tr> <tr> <td>internal</td> <td>Private IP address that load balances requests to ports 6443 and 22623 across control plane machines</td> </tr> <tr> <td>external</td> <td>Public IP address that load balances requests to port 6443 across control plane machines</td> </tr> </table> <p>If your applications create more Kubernetes LoadBalancer service objects, your cluster uses more load balancers.</p>	default	Public IP address that load balances requests to ports 80 and 443 across worker machines	internal	Private IP address that load balances requests to ports 6443 and 22623 across control plane machines	external	Public IP address that load balances requests to port 6443 across control plane machines
default	Public IP address that load balances requests to ports 80 and 443 across worker machines								
internal	Private IP address that load balances requests to ports 6443 and 22623 across control plane machines								
external	Public IP address that load balances requests to port 6443 across control plane machines								
Public IP addresses	3		Each of the two public load balancers uses a public IP address. The bootstrap machine also uses a public IP address so that you can SSH into the machine to troubleshoot issues during installation. The IP address for the bootstrap node is used only during installation.						
Private IP addresses	7		The internal load balancer, each of the three control plane machines, and each of the three worker machines each use a private IP address.						

Component	Number of components required by default	Default Azure limit	Description
Spot VM vCPUs (optional)	0 If you configure spot VMs, your cluster must have two spot VM vCPUs for every compute node.	20 per region	<p>This is an optional component. To use spot VMs, you must increase the Azure default limit to at least twice the number of compute nodes in your cluster.</p> <div style="display: flex; align-items: center;">  <div> <p>NOTE</p> <p>Using spot VMs for control plane nodes is not recommended.</p> </div> </div>

Additional resources

- [Optimizing storage](#)

7.11.3.2. Configuring a public DNS zone in Azure

To install OpenShift Container Platform, the Microsoft Azure account you use must have a dedicated public hosted DNS zone in your account. This zone must be authoritative for the domain. This service provides cluster DNS resolution and name lookup for external connections to the cluster.

Procedure

1. Identify your domain, or subdomain, and registrar. You can transfer an existing domain and registrar or obtain a new one through Azure or another source.



NOTE

For more information about purchasing domains through Azure, see [Buy a custom domain name for Azure App Service](#) in the Azure documentation.

2. If you are using an existing domain and registrar, migrate its DNS to Azure. See [Migrate an active DNS name to Azure App Service](#) in the Azure documentation.
3. Configure DNS for your domain. Follow the steps in the [Tutorial: Host your domain in Azure DNS](#) in the Azure documentation to create a public hosted zone for your domain or subdomain, extract the new authoritative name servers, and update the registrar records for the name servers that your domain uses.
Use an appropriate root domain, such as **openshiftcorp.com**, or subdomain, such as **clusters.openshiftcorp.com**.
4. If you use a subdomain, follow your company's procedures to add its delegation records to the parent domain.

You can view Azure's DNS solution by visiting this [example for creating DNS zones](#).

7.11.3.3. Increasing Azure account limits

To increase an account limit, file a support request on the Azure portal.

**NOTE**

You can increase only one type of quota per support request.

Procedure

1. From the Azure portal, click **Help + support** in the lower left corner.
2. Click **New support request** and then select the required values:
 - a. From the **Issue type** list, select **Service and subscription limits (quotas)**
 - b. From the **Subscription** list, select the subscription to modify.
 - c. From the **Quota type** list, select the quota to increase. For example, select **Compute-VM (cores-vCPUs) subscription limit increases** to increase the number of vCPUs, which is required to install a cluster.
 - d. Click **Next: Solutions**.
3. On the **Problem Details** page, provide the required information for your quota increase:
 - a. Click **Provide details** and provide the required details in the **Quota details** window.
 - b. In the SUPPORT METHOD and CONTACT INFO sections, provide the issue severity and your contact details.
4. Click **Next: Review + create** and then click **Create**.

7.11.3.4. Certificate signing requests management

Because your cluster has limited access to automatic machine management when you use infrastructure that you provision, you must provide a mechanism for approving cluster certificate signing requests (CSRs) after installation. The **kube-controller-manager** only approves the kubelet client CSRs. The **machine-approver** cannot guarantee the validity of a serving certificate that is requested by using kubelet credentials because it cannot confirm that the correct machine issued the request. You must determine and implement a method of verifying the validity of the kubelet serving certificate requests and approving them.

7.11.3.5. Recording the subscription and tenant IDs

The installation program requires the subscription and tenant IDs that are associated with your Azure account. You can use the Azure CLI to gather this information.

Prerequisites

- You have installed or updated the [Azure CLI](#).

Procedure

1. Log in to the Azure CLI by running the following command:

```
$ az login
```

2. Ensure that you are using the right subscription:

- a. View a list of available subscriptions by running the following command:

```
$ az account list --refresh
```

Example output

```
[
  {
    "cloudName": "AzureCloud",
    "id": "8xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
    "isDefault": true,
    "name": "Subscription Name 1",
    "state": "Enabled",
    "tenantId": "6xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
    "user": {
      "name": "you@example.com",
      "type": "user"
    }
  },
  {
    "cloudName": "AzureCloud",
    "id": "9xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
    "isDefault": false,
    "name": "Subscription Name 2",
    "state": "Enabled",
    "tenantId": "7xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
    "user": {
      "name": "you2@example.com",
      "type": "user"
    }
  }
]
```

- b. View the details of the active account, and confirm that this is the subscription you want to use, by running the following command:

```
$ az account show
```

Example output

```
{
  "environmentName": "AzureCloud",
  "id": "8xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
  "isDefault": true,
  "name": "Subscription Name 1",
  "state": "Enabled",
  "tenantId": "6xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
  "user": {
    "name": "you@example.com",
    "type": "user"
  }
}
```

3. If you are not using the right subscription:

- a. Change the active subscription by running the following command:

```
$ az account set -s <subscription_id>
```

- b. Verify that you are using the subscription you need by running the following command:

```
$ az account show
```

Example output

```
{
  "environmentName": "AzureCloud",
  "id": "9xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
  "isDefault": true,
  "name": "Subscription Name 2",
  "state": "Enabled",
  "tenantId": "7xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
  "user": {
    "name": "you2@example.com",
    "type": "user"
  }
}
```

4. Record the **id** and **tenantId** parameter values from the output. You require these values to install an OpenShift Container Platform cluster.

7.11.3.6. Supported identities to access Azure resources

An OpenShift Container Platform cluster requires an Azure identity to create and manage Azure resources. As such, you need one of the following types of identities to complete the installation:

- A service principal
- A system-assigned managed identity
- A user-assigned managed identity

7.11.3.7. Required Azure permissions for user-provisioned infrastructure

The installation program requires access to an Azure service principal or managed identity with the necessary permissions to deploy the cluster and to maintain its daily operation. These permissions must be granted to the Azure subscription that is associated with the identity.

The following options are available to you:

- You can assign the identity the **Contributor** and **User Access Administrator** roles. Assigning these roles is the quickest way to grant all of the required permissions. For more information about assigning roles, see the Azure documentation for [managing access to Azure resources using the Azure portal](#).
- If your organization's security policies require a more restrictive set of permissions, you can create a [custom role](#) with the necessary permissions.

The following permissions are required for creating an OpenShift Container Platform cluster on Microsoft Azure.

Example 7.64. Required permissions for creating authorization resources

- **Microsoft.Authorization/policies/audit/action**
- **Microsoft.Authorization/policies/auditIfNotExists/action**
- **Microsoft.Authorization/roleAssignments/read**
- **Microsoft.Authorization/roleAssignments/write**

Example 7.65. Required permissions for creating compute resources

- **Microsoft.Compute/images/read**
- **Microsoft.Compute/images/write**
- **Microsoft.Compute/images/delete**
- **Microsoft.Compute/availabilitySets/read**
- **Microsoft.Compute/disks/beginGetAccess/action**
- **Microsoft.Compute/disks/delete**
- **Microsoft.Compute/disks/read**
- **Microsoft.Compute/disks/write**
- **Microsoft.Compute/galleries/images/read**
- **Microsoft.Compute/galleries/images/versions/read**
- **Microsoft.Compute/galleries/images/versions/write**
- **Microsoft.Compute/galleries/images/write**
- **Microsoft.Compute/galleries/read**
- **Microsoft.Compute/galleries/write**
- **Microsoft.Compute/snapshots/read**
- **Microsoft.Compute/snapshots/write**
- **Microsoft.Compute/snapshots/delete**
- **Microsoft.Compute/virtualMachines/delete**
- **Microsoft.Compute/virtualMachines/powerOff/action**
- **Microsoft.Compute/virtualMachines/read**
- **Microsoft.Compute/virtualMachines/write**

- **Microsoft.Compute/virtualMachines/deallocate/action**

Example 7.66. Required permissions for creating identity management resources

- **Microsoft.ManagedIdentity/userAssignedIdentities/assign/action**
- **Microsoft.ManagedIdentity/userAssignedIdentities/read**
- **Microsoft.ManagedIdentity/userAssignedIdentities/write**

Example 7.67. Required permissions for creating network resources

- **Microsoft.Network/dnsZones/A/write**
- **Microsoft.Network/dnsZones/CNAME/write**
- **Microsoft.Network/dnszones/CNAME/read**
- **Microsoft.Network/dnszones/read**
- **Microsoft.Network/loadBalancers/backendAddressPools/join/action**
- **Microsoft.Network/loadBalancers/backendAddressPools/read**
- **Microsoft.Network/loadBalancers/backendAddressPools/write**
- **Microsoft.Network/loadBalancers/read**
- **Microsoft.Network/loadBalancers/write**
- **Microsoft.Network/networkInterfaces/delete**
- **Microsoft.Network/networkInterfaces/join/action**
- **Microsoft.Network/networkInterfaces/read**
- **Microsoft.Network/networkInterfaces/write**
- **Microsoft.Network/networkSecurityGroups/join/action**
- **Microsoft.Network/networkSecurityGroups/read**
- **Microsoft.Network/networkSecurityGroups/securityRules/delete**
- **Microsoft.Network/networkSecurityGroups/securityRules/read**
- **Microsoft.Network/networkSecurityGroups/securityRules/write**
- **Microsoft.Network/networkSecurityGroups/write**
- **Microsoft.Network/privateDnsZones/A/read**
- **Microsoft.Network/privateDnsZones/A/write**
- **Microsoft.Network/privateDnsZones/A/delete**

- **Microsoft.Network/privateDnsZones/SOA/read**
- **Microsoft.Network/privateDnsZones/read**
- **Microsoft.Network/privateDnsZones/virtualNetworkLinks/read**
- **Microsoft.Network/privateDnsZones/virtualNetworkLinks/write**
- **Microsoft.Network/privateDnsZones/write**
- **Microsoft.Network/publicIPAddresses/delete**
- **Microsoft.Network/publicIPAddresses/join/action**
- **Microsoft.Network/publicIPAddresses/read**
- **Microsoft.Network/publicIPAddresses/write**
- **Microsoft.Network/virtualNetworks/join/action**
- **Microsoft.Network/virtualNetworks/read**
- **Microsoft.Network/virtualNetworks/subnets/join/action**
- **Microsoft.Network/virtualNetworks/subnets/read**
- **Microsoft.Network/virtualNetworks/subnets/write**
- **Microsoft.Network/virtualNetworks/write**

Example 7.68. Required permissions for checking the health of resources

- **Microsoft.Resourcehealth/healthevent/Activated/action**
- **Microsoft.Resourcehealth/healthevent/InProgress/action**
- **Microsoft.Resourcehealth/healthevent/Pending/action**
- **Microsoft.Resourcehealth/healthevent/Resolved/action**
- **Microsoft.Resourcehealth/healthevent/Updated/action**

Example 7.69. Required permissions for creating a resource group

- **Microsoft.Resources/subscriptions/resourceGroups/read**
- **Microsoft.Resources/subscriptions/resourcegroups/write**

Example 7.70. Required permissions for creating resource tags

- **Microsoft.Resources/tags/write**

Example 7.71. Required permissions for creating storage resources

- **Microsoft.Storage/storageAccounts/blobServices/read**
- **Microsoft.Storage/storageAccounts/blobServices/containers/write**
- **Microsoft.Storage/storageAccounts/fileServices/read**
- **Microsoft.Storage/storageAccounts/fileServices/shares/read**
- **Microsoft.Storage/storageAccounts/fileServices/shares/write**
- **Microsoft.Storage/storageAccounts/fileServices/shares/delete**
- **Microsoft.Storage/storageAccounts/listKeys/action**
- **Microsoft.Storage/storageAccounts/read**
- **Microsoft.Storage/storageAccounts/write**

Example 7.72. Required permissions for creating deployments

- **Microsoft.Resources/deployments/read**
- **Microsoft.Resources/deployments/write**
- **Microsoft.Resources/deployments/validate/action**
- **Microsoft.Resources/deployments/operationstatuses/read**

Example 7.73. Optional permissions for creating compute resources

- **Microsoft.Compute/availabilitySets/delete**
- **Microsoft.Compute/availabilitySets/write**

Example 7.74. Optional permissions for creating marketplace virtual machine resources

- **Microsoft.MarketplaceOrdering/offertypes/publishers/offers/plans/agreements/read**
- **Microsoft.MarketplaceOrdering/offertypes/publishers/offers/plans/agreements/write**

Example 7.75. Optional permissions for enabling user-managed encryption

- **Microsoft.Compute/diskEncryptionSets/read**
- **Microsoft.Compute/diskEncryptionSets/write**
- **Microsoft.Compute/diskEncryptionSets/delete**
- **Microsoft.KeyVault/vaults/read**

- **Microsoft.KeyVault/vaults/write**
- **Microsoft.KeyVault/vaults/delete**
- **Microsoft.KeyVault/vaults/deploy/action**
- **Microsoft.KeyVault/vaults/keys/read**
- **Microsoft.KeyVault/vaults/keys/write**
- **Microsoft.Features/providers/features/register/action**

The following permissions are required for deleting an OpenShift Container Platform cluster on Microsoft Azure.

Example 7.76. Required permissions for deleting authorization resources

- **Microsoft.Authorization/roleAssignments/delete**

Example 7.77. Required permissions for deleting compute resources

- **Microsoft.Compute/disks/delete**
- **Microsoft.Compute/galleries/delete**
- **Microsoft.Compute/galleries/images/delete**
- **Microsoft.Compute/galleries/images/versions/delete**
- **Microsoft.Compute/virtualMachines/delete**
- **Microsoft.Compute/images/delete**

Example 7.78. Required permissions for deleting identity management resources

- **Microsoft.ManagedIdentity/userAssignedIdentities/delete**

Example 7.79. Required permissions for deleting network resources

- **Microsoft.Network/dnszones/read**
- **Microsoft.Network/dnsZones/A/read**
- **Microsoft.Network/dnsZones/A/delete**
- **Microsoft.Network/dnsZones/CNAME/read**
- **Microsoft.Network/dnsZones/CNAME/delete**
- **Microsoft.Network/loadBalancers/delete**
- **Microsoft.Network/networkInterfaces/delete**

- **Microsoft.Network/networkSecurityGroups/delete**
- **Microsoft.Network/privateDnsZones/read**
- **Microsoft.Network/privateDnsZones/A/read**
- **Microsoft.Network/privateDnsZones/delete**
- **Microsoft.Network/privateDnsZones/virtualNetworkLinks/delete**
- **Microsoft.Network/publicIPAddresses/delete**
- **Microsoft.Network/virtualNetworks/delete**

Example 7.80. Required permissions for checking the health of resources

- **Microsoft.Resourcehealth/healthevent/Activated/action**
- **Microsoft.Resourcehealth/healthevent/Resolved/action**
- **Microsoft.Resourcehealth/healthevent/Updated/action**

Example 7.81. Required permissions for deleting a resource group

- **Microsoft.Resources/subscriptions/resourcegroups/delete**

Example 7.82. Required permissions for deleting storage resources

- **Microsoft.Storage/storageAccounts/delete**
- **Microsoft.Storage/storageAccounts/listKeys/action**



NOTE

To install OpenShift Container Platform on Azure, you must scope the permissions related to resource group creation to your subscription. After the resource group is created, you can scope the rest of the permissions to the created resource group. If the public DNS zone is present in a different resource group, then the network DNS zone related permissions must always be applied to your subscription.

You can scope all the permissions to your subscription when deleting an OpenShift Container Platform cluster.

7.11.3.8. Using Azure managed identities

The installation program requires an Azure identity to complete the installation. You can use either a system-assigned or user-assigned managed identity.

If you are unable to use a managed identity, you can use a service principal.

Procedure

1. If you are using a system-assigned managed identity, enable it on the virtual machine that you will run the installation program from.
2. If you are using a user-assigned managed identity:
 - a. Assign it to the virtual machine that you will run the installation program from.
 - b. Record its client ID. You require this value when installing the cluster.
For more information about viewing the details of a user-assigned managed identity, see the Microsoft Azure documentation for [listing user-assigned managed identities](#).
3. Verify that the required permissions are assigned to the managed identity.

7.11.3.9. Creating a service principal

The installation program requires an Azure identity to complete the installation. You can use a service principal.

If you are unable to use a service principal, you can use a managed identity.

Prerequisites

- You have installed or updated the [Azure CLI](#).
- You have an Azure subscription ID.
- If you are not going to assign the **Contributor** and **User Administrator Access** roles to the service principal, you have created a custom role with the required Azure permissions.

Procedure

1. Create the service principal for your account by running the following command:

```
$ az ad sp create-for-rbac --role <role_name> \ 1
--name <service_principal> \ 2
--scopes /subscriptions/<subscription_id> 3
```

- 1 Defines the role name. You can use the **Contributor** role, or you can specify a custom role which contains the necessary permissions.
- 2 Defines the service principal name.
- 3 Specifies the subscription ID.

Example output

```
Creating 'Contributor' role assignment under scope '/subscriptions/<subscription_id>'
The output includes credentials that you must protect. Be sure that you do not
include these credentials in your code or check the credentials into your source
control. For more information, see https://aka.ms/azadsp-cli
{
  "appId": "axxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
  "displayName": <service_principal>,

```

```
"password": "00000000-0000-0000-0000-000000000000",  
"tenantId": "8xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxxx"  
}
```

- Record the values of the **appld** and **password** parameters from the output. You require these values when installing the cluster.
- If you applied the **Contributor** role to your service principal, assign the **User Administrator Access** role by running the following command:

```
$ az role assignment create --role "User Access Administrator" \  
--assignee-object-id $(az ad sp show --id <appld> --query id -o tsv) 1  
--scope /subscriptions/<subscription_id> 2
```

- 1** Specify the **appld** parameter value for your service principal.
- 2** Specifies the subscription ID.

Additional resources

- For more information about CCO modes, see [About the Cloud Credential Operator](#).

7.11.3.10. Supported Azure regions

The installation program dynamically generates the list of available Microsoft Azure regions based on your subscription.

Supported Azure public regions

- **australiacentral** (Australia Central)
- **australiaeast** (Australia East)
- **australiasoutheast** (Australia South East)
- **brazilsouth** (Brazil South)
- **canadacentral** (Canada Central)
- **canadaeast** (Canada East)
- **centralindia** (Central India)
- **centralus** (Central US)
- **eastasia** (East Asia)
- **eastus** (East US)
- **eastus2** (East US 2)
- **francecentral** (France Central)
- **germanywestcentral** (Germany West Central)

- **israelcentral** (Israel Central)
- **italynorth** (Italy North)
- **japaneast** (Japan East)
- **japanwest** (Japan West)
- **koreacentral** (Korea Central)
- **koreasouth** (Korea South)
- **mexicocentral** (Mexico Central)
- **northcentralus** (North Central US)
- **northeurope** (North Europe)
- **norwayeast** (Norway East)
- **polandcentral** (Poland Central)
- **qatarcentral** (Qatar Central)
- **southafricanorth** (South Africa North)
- **southcentralus** (South Central US)
- **southeastasia** (Southeast Asia)
- **southindia** (South India)
- **swedencentral** (Sweden Central)
- **switzerlandnorth** (Switzerland North)
- **uaenorth** (UAE North)
- **uksouth** (UK South)
- **ukwest** (UK West)
- **westcentralus** (West Central US)
- **westeurope** (West Europe)
- **westindia** (West India)
- **westus** (West US)
- **westus2** (West US 2)
- **westus3** (West US 3)

Supported Azure Government regions

Support for the following Microsoft Azure Government (MAG) regions was added in OpenShift Container Platform version 4.6:

- **usgovtexas** (US Gov Texas)
- **usgovvirginia** (US Gov Virginia)

You can reference all available MAG regions in the [Azure documentation](#). Other provided MAG regions are expected to work with OpenShift Container Platform, but have not been tested.

7.11.4. Requirements for a cluster with user-provisioned infrastructure

For a cluster that contains user-provisioned infrastructure, you must deploy all of the required machines.


This section describes the requirements for deploying OpenShift Container Platform on user-provisioned infrastructure.

7.11.4.1. Required machines for cluster installation

The smallest OpenShift Container Platform clusters require the following hosts:

Table 7.25. Minimum required hosts

Hosts	Description
One temporary bootstrap machine	The cluster requires the bootstrap machine to deploy the OpenShift Container Platform cluster on the three control plane machines. You can remove the bootstrap machine after you install the cluster.
Three control plane machines	The control plane machines run the Kubernetes and OpenShift Container Platform services that form the control plane.
At least two compute machines, which are also known as worker machines.	The workloads requested by OpenShift Container Platform users run on the compute machines.



IMPORTANT

To maintain high availability of your cluster, use separate physical hosts for these cluster machines.

The bootstrap and control plane machines must use Red Hat Enterprise Linux CoreOS (RHCOS) as the operating system. However, the compute machines can choose between Red Hat Enterprise Linux CoreOS (RHCOS), Red Hat Enterprise Linux (RHEL) 8.6 and later.

Note that RHCOS is based on Red Hat Enterprise Linux (RHEL) 9.2 and inherits all of its hardware certifications and requirements. See [Red Hat Enterprise Linux technology capabilities and limits](#) .

7.11.4.2. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 7.26. Minimum resource requirements

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Bootstrap	RHCOS	4	16 GB	100 GB	300
Control plane	RHCOS	4	16 GB	100 GB	300
Compute	RHCOS, RHEL 8.6 and later [3]	2	8 GB	100 GB	300

1. One vCPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio: (threads per core × cores) × sockets = vCPUs.
2. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes which require a 10 ms p99 fsync duration. Note that on many cloud platforms, storage size and IOPS scale together, so you might need to over-allocate storage volume to obtain sufficient performance.
3. As with all user-provisioned installations, if you choose to use RHEL compute machines in your cluster, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Use of RHEL 7 compute machines is deprecated and has been removed in OpenShift Container Platform 4.10 and later.



NOTE

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).



IMPORTANT

You are required to use Azure virtual machines that have the **premiumIO** parameter set to **true**.

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

7.11.4.3. Tested instance types for Azure

The following Microsoft Azure instance types have been tested with OpenShift Container Platform.

Example 7.83. Machine types based on 64-bit x86 architecture

- **c4.***
- **c5.***
- **c5a.***
- **i3.***
- **m4.***
- **m5.***
- **m5a.***
- **m6a.***
- **m6i.***
- **r4.***
- **r5.***
- **r5a.***
- **r6i.***
- **t3.***
- **t3a.***

7.11.4.4. Tested instance types for Azure on 64-bit ARM infrastructures

The following Microsoft Azure ARM64 instance types have been tested with OpenShift Container Platform.

Example 7.84. Machine types based on 64-bit ARM architecture

- **c6g.***
- **m6g.***

7.11.5. Using the Azure Marketplace offering

Using the Azure Marketplace offering lets you deploy an OpenShift Container Platform cluster, which is billed on pay-per-use basis (hourly, per core) through Azure, while still being supported directly by Red Hat.

To deploy an OpenShift Container Platform cluster using the Azure Marketplace offering, you must first obtain the Azure Marketplace image. The installation program uses this image to deploy worker or control plane nodes. When obtaining your image, consider the following:

- While the images are the same, the Azure Marketplace publisher is different depending on your region. If you are located in North America, specify **redhat** as the publisher. If you are located in EMEA, specify **redhat-limited** as the publisher.
- The offer includes a **rh-ocp-worker** SKU and a **rh-ocp-worker-gen1** SKU. The **rh-ocp-worker** SKU represents a Hyper-V generation version 2 VM image. The default instance types used in OpenShift Container Platform are version 2 compatible. If you plan to use an instance type that is only version 1 compatible, use the image associated with the **rh-ocp-worker-gen1** SKU. The **rh-ocp-worker-gen1** SKU represents a Hyper-V version 1 VM image.



IMPORTANT

Installing images with the Azure marketplace is not supported on clusters with 64-bit ARM instances.

Prerequisites

- You have installed the Azure CLI client (**az**).
- Your Azure account is entitled for the offer and you have logged into this account with the Azure CLI client.

Procedure

1. Display all of the available OpenShift Container Platform images by running one of the following commands:

- North America:

```
$ az vm image list --all --offer rh-ocp-worker --publisher redhat -o table
```

Example output

Offer	Publisher	Sku	Urn	Version
rh-ocp-worker	RedHat	rh-ocp-worker	RedHat:rh-ocp-worker:rh-ocp-worker:413.92.2023101700	413.92.2023101700
rh-ocp-worker-gen1	RedHat	rh-ocp-worker-gen1	RedHat:rh-ocp-worker:rh-ocp-worker-gen1:413.92.2023101700	413.92.2023101700

- EMEA:

```
$ az vm image list --all --offer rh-ocp-worker --publisher redhat-limited -o table
```

Example output

Offer	Publisher	Sku	Urn	Version
-------	-----------	-----	-----	---------

```

-----
rh-ocp-worker redhat-limited rh-ocp-worker redhat-limited:rh-ocp-worker:rh-ocp-
worker:413.92.2023101700 413.92.2023101700
rh-ocp-worker redhat-limited rh-ocp-worker-gen1 redhat-limited:rh-ocp-worker:rh-ocp-
worker-gen1:413.92.2023101700 413.92.2023101700

```

**NOTE**

Use the latest image that is available for compute and control plane nodes. If required, your VMs are automatically upgraded as part of the installation process.

2. Inspect the image for your offer by running one of the following commands:

- North America:

```
$ az vm image show --urn redhat:rh-ocp-worker:rh-ocp-worker:<version>
```

- EMEA:

```
$ az vm image show --urn redhat-limited:rh-ocp-worker:rh-ocp-worker:<version>
```

3. Review the terms of the offer by running one of the following commands:

- North America:

```
$ az vm image terms show --urn redhat:rh-ocp-worker:rh-ocp-worker:<version>
```

- EMEA:

```
$ az vm image terms show --urn redhat-limited:rh-ocp-worker:rh-ocp-worker:<version>
```

4. Accept the terms of the offering by running one of the following commands:

- North America:

```
$ az vm image terms accept --urn redhat:rh-ocp-worker:rh-ocp-worker:<version>
```

- EMEA:

```
$ az vm image terms accept --urn redhat-limited:rh-ocp-worker:rh-ocp-worker:<version>
```

5. Record the image details of your offer. If you use the Azure Resource Manager (ARM) template to deploy your compute nodes:

- Update **storageProfile.imageReference** by deleting the **id** parameter and adding the **offer**, **publisher**, **sku**, and **version** parameters by using the values from your offer.
- Specify a **plan** for the virtual machines (VMs).

Example 06_workers.json ARM template with an updated storageProfile.imageReference object and a specified plan

```
...
```



```

"plan" : {
  "name": "rh-ocp-worker",
  "product": "rh-ocp-worker",
  "publisher": "redhat"
},
"dependsOn" : [
  "[concat('Microsoft.Network/networkInterfaces/', concat(variables('vmNames')
[copyIndex()], '-nic'))]"
],
"properties" : {
...
"storageProfile": {
  "imageReference": {
    "offer": "rh-ocp-worker",
    "publisher": "redhat",
    "sku": "rh-ocp-worker",
    "version": "413.92.2023101700"
  }
...
}
...
}

```

7.11.6. Obtaining the installation program


Before you install OpenShift Container Platform, download the installation file on the host you are using for installation.

Prerequisites

- You have a computer that runs Linux or macOS, with 500 MB of local disk space.


Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select your infrastructure provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.



IMPORTANT

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.



IMPORTANT

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.

4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#) . This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

7.11.7. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.



NOTE

You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1** Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.

**NOTE**

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.

**NOTE**

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```

**NOTE**

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program. If you install a cluster on infrastructure that you provision, you must provide the key to the installation program.

7.11.8. Creating the installation files for Azure

To install OpenShift Container Platform on Microsoft Azure using user-provisioned infrastructure, you must generate the files that the installation program needs to deploy your cluster and modify them so that the cluster creates only the machines that it will use. You generate and customize the **install-config.yaml** file, Kubernetes manifests, and Ignition config files. You also have the option to first set up a separate **var** partition during the preparation phases of installation.

7.11.8.1. Optional: Creating a separate /var partition

It is recommended that disk partitioning for OpenShift Container Platform be left to the installer. However, there are cases where you might want to create separate partitions in a part of the filesystem that you expect to grow.

OpenShift Container Platform supports the addition of a single partition to attach storage to either the **/var** partition or a subdirectory of **/var**. For example:

- **/var/lib/containers**: Holds container-related content that can grow as more images and containers are added to a system.
- **/var/lib/etcd**: Holds data that you might want to keep separate for purposes such as performance optimization of etcd storage.
- **/var**: Holds data that you might want to keep separate for purposes such as auditing.

Storing the contents of a **/var** directory separately makes it easier to grow storage for those areas as needed and reinstall OpenShift Container Platform at a later date and keep that data intact. With this method, you will not have to pull all your containers again, nor will you have to copy massive log files when you update systems.

Because **/var** must be in place before a fresh installation of Red Hat Enterprise Linux CoreOS (RHCOS), the following procedure sets up the separate **/var** partition by creating a machine config manifest that is inserted during the **openshift-install** preparation phases of an OpenShift Container Platform installation.



IMPORTANT

If you follow the steps to create a separate **/var** partition in this procedure, it is not necessary to create the Kubernetes manifest and Ignition config files again as described later in this section.

Procedure

1. Create a directory to hold the OpenShift Container Platform installation files:

```
$ mkdir $HOME/clusterconfig
```

2. Run **openshift-install** to create a set of files in the **manifest** and **openshift** subdirectories. Answer the system questions as you are prompted:

```
$ openshift-install create manifests --dir $HOME/clusterconfig
```

-

Example output

```
? SSH Public Key ...
INFO Credentials loaded from the "myprofile" profile in file "/home/myuser/.aws/credentials"
INFO Consuming Install Config from target directory
INFO Manifests created in: $HOME/clusterconfig/manifests and
$HOME/clusterconfig/openshift
```

- Optional: Confirm that the installation program created manifests in the **clusterconfig/openshift** directory:

```
$ ls $HOME/clusterconfig/openshift/
```

Example output

```
99_kubeadmin-password-secret.yaml
99_openshift-cluster-api_master-machines-0.yaml
99_openshift-cluster-api_master-machines-1.yaml
99_openshift-cluster-api_master-machines-2.yaml
...
```

- Create a Butane config that configures the additional partition. For example, name the file **\$HOME/clusterconfig/98-var-partition.bu**, change the disk device name to the name of the storage device on the **worker** systems, and set the storage size as appropriate. This example places the **/var** directory on a separate partition:

```
variant: openshift
version: 4.16.0
metadata:
  labels:
    machineconfiguration.openshift.io/role: worker
  name: 98-var-partition
storage:
  disks:
    - device: /dev/disk/by-id/<device_name> ❶
      partitions:
        - label: var
          start_mib: <partition_start_offset> ❷
          size_mib: <partition_size> ❸
          number: 5
      filesystems:
        - device: /dev/disk/by-partlabel/var
          path: /var
          format: xfs
          mount_options: [defaults, prjquota] ❹
          with_mount_unit: true
```

- The storage device name of the disk that you want to partition.
- When adding a data partition to the boot disk, a minimum value of 25000 MiB (Mebibytes) is recommended. The root file system is automatically resized to fill all available space up to the specified offset. If no value is specified, or if the specified value is smaller than the recommended minimum, the resulting root file system will be too small, and future

reinstalls of RHCOS might overwrite the beginning of the data partition.

- 3 The size of the data partition in mebibytes.
- 4 The **prjquota** mount option must be enabled for filesystems used for container storage.



NOTE

When creating a separate **/var** partition, you cannot use different instance types for worker nodes, if the different instance types do not have the same device name.

5. Create a manifest from the Butane config and save it to the **clusterconfig/openshift** directory. For example, run the following command:

```
$ butane $HOME/clusterconfig/98-var-partition.bu -o $HOME/clusterconfig/openshift/98-var-partition.yaml
```

6. Run **openshift-install** again to create Ignition configs from a set of files in the **manifest** and **openshift** subdirectories:

```
$ openshift-install create ignition-configs --dir $HOME/clusterconfig
$ ls $HOME/clusterconfig/
auth bootstrap.ign master.ign metadata.json worker.ign
```

Now you can use the Ignition config files as input to the installation procedures to install Red Hat Enterprise Linux CoreOS (RHCOS) systems.

7.11.8.2. Creating the installation configuration file

You can customize the OpenShift Container Platform cluster you install on Microsoft Azure.

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have an Azure subscription ID and tenant ID.
- If you are installing the cluster using a service principal, you have its application ID and password.
- If you are installing the cluster using a system-assigned managed identity, you have enabled it on the virtual machine that you will run the installation program from.
- If you are installing the cluster using a user-assigned managed identity, you have met these prerequisites:
 - You have its client ID.
 - You have assigned it to the virtual machine that you will run the installation program from.

Procedure

1. Optional: If you have run the installation program on this computer before, and want to use an alternative service principal or managed identity, go to the `~/.azure/` directory and delete the **osServicePrincipal.json** configuration file.
Deleting this file prevents the installation program from automatically reusing subscription and authentication values from a previous installation.
2. Create the **install-config.yaml** file.

- a. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install create install-config --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the directory name to store the files that the installation program creates.

When specifying the directory:

- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
- Use an empty directory. Some installation assets, such as bootstrap X.509 certificates, have short expiration intervals, therefore you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

- b. At the prompts, provide the configuration details for your cloud:
 - i. Optional: Select an SSH key to use to access your cluster machines.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- ii. Select **azure** as the platform to target.
If the installation program cannot locate the **osServicePrincipal.json** configuration file from a previous installation, you are prompted for Azure subscription and authentication values.
- iii. Enter the following Azure parameter values for your subscription:
 - **azure subscription id** Enter the subscription ID to use for the cluster.
 - **azure tenant id** Enter the tenant ID.
- iv. Depending on the Azure identity you are using to deploy the cluster, do one of the following when prompted for the **azure service principal client id**
 - If you are using a service principal, enter its application ID.
 - If you are using a system-assigned managed identity, leave this value blank.

- If you are using a user-assigned managed identity, specify its client ID.
- v. Depending on the Azure identity you are using to deploy the cluster, do one of the following when prompted for the **azure service principal client secret**
 - If you are using a service principal, enter its password.
 - If you are using a system-assigned managed identity, leave this value blank.
 - If you are using a user-assigned managed identity, leave this value blank.
 - vi. Select the region to deploy the cluster to.
 - vii. Select the base domain to deploy the cluster to. The base domain corresponds to the Azure DNS Zone that you created for your cluster.
 - viii. Enter a descriptive name for your cluster.



IMPORTANT

All Azure resources that are available through public endpoints are subject to resource name restrictions, and you cannot create resources that use certain terms. For a list of terms that Azure restricts, see [Resolve reserved resource name errors](#) in the Azure documentation.

3. Modify the **install-config.yaml** file. You can find more information about the available parameters in the "Installation configuration parameters" section.



NOTE

If you are installing a three-node cluster, be sure to set the **compute.replicas** parameter to **0**. This ensures that the cluster's control planes are schedulable. For more information, see "Installing a three-node cluster on Azure".

4. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

If previously not detected, the installation program creates an **osServicePrincipal.json** configuration file and stores this file in the `~/.azure/` directory on your computer. This ensures that the installation program can load the profile when it is creating an OpenShift Container Platform cluster on the target platform.

7.11.8.3. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when

http/https proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.



NOTE

The installation program does not support the proxy **readinessEndpoints** field.



NOTE

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$. /openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.



NOTE

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

7.11.8.4. Exporting common variables for ARM templates

You must export a common set of variables that are used with the provided Azure Resource Manager (ARM) templates used to assist in completing a user-provided infrastructure install on Microsoft Azure.



NOTE

Specific ARM templates can also require additional exported variables, which are detailed in their related procedures.

Prerequisites

- Obtain the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Export common variables found in the **install-config.yaml** to be used by the provided ARM templates:

```
$ export CLUSTER_NAME=<cluster_name> 1
$ export AZURE_REGION=<azure_region> 2
$ export SSH_KEY=<ssh_key> 3
$ export BASE_DOMAIN=<base_domain> 4
$ export BASE_DOMAIN_RESOURCE_GROUP=<base_domain_resource_group> 5
```

- 1 The value of the **.metadata.name** attribute from the **install-config.yaml** file.
- 2 The region to deploy the cluster into, for example **centralus**. This is the value of the **.platform.azure.region** attribute from the **install-config.yaml** file.
- 3 The SSH RSA public key file as a string. You must enclose the SSH key in quotes since it contains spaces. This is the value of the **.sshKey** attribute from the **install-config.yaml** file.
- 4 The base domain to deploy the cluster to. The base domain corresponds to the public DNS zone that you created for your cluster. This is the value of the **.baseDomain** attribute from the **install-config.yaml** file.
- 5 The resource group where the public DNS zone exists. This is the value of the **.platform.azure.baseDomainResourceGroupName** attribute from the **install-config.yaml** file.

For example:

```
$ export CLUSTER_NAME=test-cluster
$ export AZURE_REGION=centralus
$ export SSH_KEY="ssh-rsa xxx/xxx/xxx= user@email.com"
$ export BASE_DOMAIN=example.com
$ export BASE_DOMAIN_RESOURCE_GROUP=ocp-cluster
```

2. Export the kubeadmin credentials:


```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

7.11.8.5. Creating the Kubernetes manifest and Ignition config files

Because you must modify some cluster definition files and manually start the cluster machines, you must generate the Kubernetes manifest and Ignition config files that the cluster needs to configure the machines.

The installation configuration file transforms into the Kubernetes manifests. The manifests wrap into the Ignition configuration files, which are later used to configure the cluster machines.



IMPORTANT

- The Ignition config files that the OpenShift Container Platform installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

Prerequisites

- You obtained the OpenShift Container Platform installation program.
- You created the **install-config.yaml** installation configuration file.

Procedure

1. Change to the directory that contains the OpenShift Container Platform installation program and generate the Kubernetes manifests for the cluster:

```
$ ./openshift-install create manifests --dir <installation_directory> 1
```

- 1 For **<installation_directory>**, specify the installation directory that contains the **install-config.yaml** file you created.

2. Remove the Kubernetes manifest files that define the control plane machines:

```
$ rm -f <installation_directory>/openshift/99_openshift-cluster-api_master-machines-*.yaml
```


By removing these files, you prevent the cluster from automatically generating control plane machines.

3. Remove the Kubernetes manifest files that define the control plane machine set:

```
$ rm -f <installation_directory>/openshift/99_openshift-machine-api_master-control-plane-machine-set.yaml
```

4. Remove the Kubernetes manifest files that define the worker machines:

```
$ rm -f <installation_directory>/openshift/99_openshift-cluster-api_worker-machineset-*.yaml
```



IMPORTANT


If you disabled the **MachineAPI** capability when installing a cluster on user-provisioned infrastructure, you must remove the Kubernetes manifest files that define the worker machines. Otherwise, your cluster fails to install.

Because you create and manage the worker machines yourself, you do not need to initialize these machines.



WARNING

If you are installing a three-node cluster, skip the following step to allow the control plane nodes to be schedulable.



IMPORTANT

When you configure control plane nodes from the default unschedulable to schedulable, additional subscriptions are required. This is because control plane nodes then become compute nodes.

5. Check that the **mastersSchedulable** parameter in the `<installation_directory>/manifests/cluster-scheduler-02-config.yml` Kubernetes manifest file is set to **false**. This setting prevents pods from being scheduled on the control plane machines:
 - a. Open the `<installation_directory>/manifests/cluster-scheduler-02-config.yml` file.
 - b. Locate the **mastersSchedulable** parameter and ensure that it is set to **false**.
 - c. Save and exit the file.
6. Optional: If you do not want [the Ingress Operator](#) to create DNS records on your behalf, remove the **privateZone** and **publicZone** sections from the `<installation_directory>/manifests/cluster-dns-02-config.yml` DNS configuration file:

```
apiVersion: config.openshift.io/v1
kind: DNS
metadata:
  creationTimestamp: null
  name: cluster
spec:
  baseDomain: example.openshift.com
  privateZone: ❶
    id: mycluster-100419-private-zone
  publicZone: ❷
    id: example.openshift.com
status: {}
```

- ❶ ❷ Remove this section completely.

If you do so, you must add ingress DNS records manually in a later step.

7. When configuring Azure on user-provisioned infrastructure, you must export some common variables defined in the manifest files to use later in the Azure Resource Manager (ARM) templates:

- a. Export the infrastructure ID by using the following command:

- a. Export the infrastructure ID by using the following command:

```
$ export INFRA_ID=<infra_id> 1
```

- 1 The OpenShift Container Platform cluster has been assigned an identifier (**INFRA_ID**) in the form of **<cluster_name>-<random_string>**. This will be used as the base name for most resources created using the provided ARM templates. This is the value of the **.status.infrastructureName** attribute from the **manifests/cluster-infrastructure-02-config.yml** file.

- b. Export the resource group by using the following command:

```
$ export RESOURCE_GROUP=<resource_group> 1
```

- 1 All resources created in this Azure deployment exists as part of a [resource group](#). The resource group name is also based on the **INFRA_ID**, in the form of **<cluster_name>-<random_string>-rg**. This is the value of the **.status.platformStatus.azure.resourceGroupName** attribute from the **manifests/cluster-infrastructure-02-config.yml** file.

8. To create the Ignition configuration files, run the following command from the directory that contains the installation program:

```
$. /openshift-install create ignition-configs --dir <installation_directory> 1
```

- 1 For **<installation_directory>**, specify the same installation directory.

Ignition config files are created for the bootstrap, control plane, and compute nodes in the installation directory. The **kubeadmin-password** and **kubeconfig** files are created in the **./<installation_directory>/auth** directory:

```
.
├── auth
│   ├── kubeadmin-password
│   └── kubeconfig
├── bootstrap.ign
├── master.ign
├── metadata.json
└── worker.ign
```

7.11.9. Creating the Azure resource group

You must create a Microsoft Azure [resource group](#) and an identity for that resource group. These are both used during the installation of your OpenShift Container Platform cluster on Azure.

Prerequisites

- Configure an Azure account.
- Generate the Ignition config files for your cluster.

Procedure

1. Create the resource group in a supported Azure region:

```
$ az group create --name ${RESOURCE_GROUP} --location ${AZURE_REGION}
```

2. Create an Azure identity for the resource group:

```
$ az identity create -g ${RESOURCE_GROUP} -n ${INFRA_ID}-identity
```

This is used to grant the required access to Operators in your cluster. For example, this allows the Ingress Operator to create a public IP and its load balancer. You must assign the Azure identity to a role.

3. Grant the Contributor role to the Azure identity:

- a. Export the following variables required by the Azure role assignment:

```
$ export PRINCIPAL_ID=`az identity show -g ${RESOURCE_GROUP} -n ${INFRA_ID}-identity --query principalId --out tsv`
```

```
$ export RESOURCE_GROUP_ID=`az group show -g ${RESOURCE_GROUP} --query id --out tsv`
```

- b. Assign the Contributor role to the identity:

```
$ az role assignment create --assignee "${PRINCIPAL_ID}" --role 'Contributor' --scope "${RESOURCE_GROUP_ID}"
```



NOTE

If you want to assign a custom role with all the required permissions to the identity, run the following command:

```
$ az role assignment create --assignee "${PRINCIPAL_ID}" --role <custom_role> \ 1 --scope "${RESOURCE_GROUP_ID}"
```

- 1** Specifies the custom role name.

7.11.10. Uploading the RHCOS cluster image and bootstrap Ignition config file

The Azure client does not support deployments based on files existing locally. You must copy and store the RHCOS virtual hard disk (VHD) cluster image and bootstrap Ignition config file in a storage container so they are accessible during deployment.

Prerequisites

- Configure an Azure account.
- Generate the Ignition config files for your cluster.

Procedure

1. Create an Azure storage account to store the VHD cluster image:

```
$ az storage account create -g ${RESOURCE_GROUP} --location ${AZURE_REGION} --name ${CLUSTER_NAME}sa --kind Storage --sku Standard_LRS
```



WARNING

The Azure storage account name must be between 3 and 24 characters in length and use numbers and lower-case letters only. If your **CLUSTER_NAME** variable does not follow these restrictions, you must manually define the Azure storage account name. For more information on Azure storage account name restrictions, see [Resolve errors for storage account names](#) in the Azure documentation.

2. Export the storage account key as an environment variable:

```
$ export ACCOUNT_KEY=`az storage account keys list -g ${RESOURCE_GROUP} --account-name ${CLUSTER_NAME}sa --query "[0].value" -o tsv`
```

3. Export the URL of the RHCOS VHD to an environment variable:

```
$ export VHD_URL=`openshift-install coreos print-stream-json | jq -r '.architectures.<architecture>."rhel-coreos-extensions"."azure-disk".url`
```

where:

<architecture>

Specifies the architecture, valid values include **x86_64** or **aarch64**.



IMPORTANT

The RHCOS images might not change with every release of OpenShift Container Platform. You must specify an image with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Use the image version that matches your OpenShift Container Platform version if it is available.

4. Create the storage container for the VHD:

```
$ az storage container create --name vhd --account-name ${CLUSTER_NAME}sa --account-key ${ACCOUNT_KEY}
```

5. Copy the local VHD to a blob:

```
$ az storage blob copy start --account-name ${CLUSTER_NAME}sa --account-key ${ACCOUNT_KEY} --destination-blob "rhcos.vhd" --destination-container vhd --source-uri "${VHD_URL}"
```


6. Create a blob storage container and upload the generated **bootstrap.ign** file:


```
$ az storage container create --name files --account-name ${CLUSTER_NAME}sa --
account-key ${ACCOUNT_KEY}
```

```
$ az storage blob upload --account-name ${CLUSTER_NAME}sa --account-key
${ACCOUNT_KEY} -c "files" -f "<installation_directory>/bootstrap.ign" -n "bootstrap.ign"
```

7.11.11. Example for creating DNS zones

DNS records are required for clusters that use user-provisioned infrastructure. You should choose the DNS strategy that fits your scenario.

For this example, [Azure's DNS solution](#) is used, so you will create a new public DNS zone for external (internet) visibility and a private DNS zone for internal cluster resolution.



NOTE

The public DNS zone is not required to exist in the same resource group as the cluster deployment and might already exist in your organization for the desired base domain. If that is the case, you can skip creating the public DNS zone; be sure the installation config you generated earlier reflects that scenario.

Prerequisites

- Configure an Azure account.
- Generate the Ignition config files for your cluster.

Procedure

1. Create the new public DNS zone in the resource group exported in the **BASE_DOMAIN_RESOURCE_GROUP** environment variable:

```
$ az network dns zone create -g ${BASE_DOMAIN_RESOURCE_GROUP} -n
${CLUSTER_NAME}.${BASE_DOMAIN}
```

You can skip this step if you are using a public DNS zone that already exists.

2. Create the private DNS zone in the same resource group as the rest of this deployment:

```
$ az network private-dns zone create -g ${RESOURCE_GROUP} -n
${CLUSTER_NAME}.${BASE_DOMAIN}
```

You can learn more about [configuring a public DNS zone in Azure](#) by visiting that section.

7.11.12. Creating a VNet in Azure

You must create a virtual network (VNet) in Microsoft Azure for your OpenShift Container Platform cluster to use. You can customize the VNet to meet your requirements. One way to create the VNet is to modify the provided Azure Resource Manager (ARM) template.

**NOTE**

If you do not use the provided ARM template to create your Azure infrastructure, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- Configure an Azure account.
- Generate the Ignition config files for your cluster.

Procedure

1. Copy the template from the **ARM template for the VNet** section of this topic and save it as **01_vnet.json** in your cluster's installation directory. This template describes the VNet that your cluster requires.
2. Create the deployment by using the **az** CLI:

```
$ az deployment group create -g ${RESOURCE_GROUP} \
  --template-file "<installation_directory>/01_vnet.json" \
  --parameters baseName="${INFRA_ID}" 1
```

- 1** The base name to be used in resource names; this is usually the cluster's infrastructure ID.

3. Link the VNet template to the private DNS zone:

```
$ az network private-dns link vnet create -g ${RESOURCE_GROUP} -z
  ${CLUSTER_NAME}.${BASE_DOMAIN} -n ${INFRA_ID}-network-link -v "${INFRA_ID}-vnet"
  -e false
```

7.11.12.1. ARM template for the VNet

You can use the following Azure Resource Manager (ARM) template to deploy the VNet that you need for your OpenShift Container Platform cluster:

Example 7.85.01_vnet.json ARM template

```
{
  "$schema" : "https://schema.management.azure.com/schemas/2015-01-01/deploymentTemplate.json#",
  "contentVersion" : "1.0.0.0",
  "parameters" : {
    "baseName" : {
      "type" : "string",
      "minLength" : 1,
      "metadata" : {
        "description" : "Base name to be used in resource names (usually the cluster's Infra ID)"
      }
    }
  }
},
```

```

"variables" : {
  "location" : "[resourceGroup().location]",
  "virtualNetworkName" : "[concat(parameters('baseName'), '-vnet')]",
  "addressPrefix" : "10.0.0.0/16",
  "masterSubnetName" : "[concat(parameters('baseName'), '-master-subnet')]",
  "masterSubnetPrefix" : "10.0.0.0/24",
  "nodeSubnetName" : "[concat(parameters('baseName'), '-worker-subnet')]",
  "nodeSubnetPrefix" : "10.0.1.0/24",
  "clusterNsgName" : "[concat(parameters('baseName'), '-nsg')]"
},
"resources" : [
  {
    "apiVersion" : "2018-12-01",
    "type" : "Microsoft.Network/virtualNetworks",
    "name" : "[variables('virtualNetworkName')]",
    "location" : "[variables('location')]",
    "dependsOn" : [
      "[concat('Microsoft.Network/networkSecurityGroups/', variables('clusterNsgName'))]"
    ],
    "properties" : {
      "addressSpace" : {
        "addressPrefixes" : [
          "[variables('addressPrefix')]"
        ]
      },
      "subnets" : [
        {
          "name" : "[variables('masterSubnetName')]",
          "properties" : {
            "addressPrefix" : "[variables('masterSubnetPrefix')]",
            "serviceEndpoints": [],
            "networkSecurityGroup" : {
              "id" : "[resourceId('Microsoft.Network/networkSecurityGroups',
variables('clusterNsgName'))]"
            }
          }
        },
        {
          "name" : "[variables('nodeSubnetName')]",
          "properties" : {
            "addressPrefix" : "[variables('nodeSubnetPrefix')]",
            "serviceEndpoints": [],
            "networkSecurityGroup" : {
              "id" : "[resourceId('Microsoft.Network/networkSecurityGroups',
variables('clusterNsgName'))]"
            }
          }
        }
      ]
    }
  },
  {
    "type" : "Microsoft.Network/networkSecurityGroups",
    "name" : "[variables('clusterNsgName')]",
    "apiVersion" : "2018-10-01",
    "location" : "[variables('location')]",

```

```

"properties" : {
  "securityRules" : [
    {
      "name" : "apiserver_in",
      "properties" : {
        "protocol" : "Tcp",
        "sourcePortRange" : "*",
        "destinationPortRange" : "6443",
        "sourceAddressPrefix" : "*",
        "destinationAddressPrefix" : "*",
        "access" : "Allow",
        "priority" : 101,
        "direction" : "Inbound"
      }
    }
  ]
}

```

7.11.13. Deploying the RHCOS cluster image for the Azure infrastructure

You must use a valid Red Hat Enterprise Linux CoreOS (RHCOS) image for Microsoft Azure for your OpenShift Container Platform nodes.

Prerequisites

- Configure an Azure account.
- Generate the Ignition config files for your cluster.
- Store the RHCOS virtual hard disk (VHD) cluster image in an Azure storage container.
- Store the bootstrap Ignition config file in an Azure storage container.

Procedure

1. Copy the template from the **ARM template for image storage** section of this topic and save it as **02_storage.json** in your cluster's installation directory. This template describes the image storage that your cluster requires.
2. Export the RHCOS VHD blob URL as a variable:

```
$ export VHD_BLOB_URL=`az storage blob url --account-name ${CLUSTER_NAME}sa --
account-key ${ACCOUNT_KEY} -c vhd -n "rhcos.vhd" -o tsv`
```

3. Deploy the cluster image:

```
$ az deployment group create -g ${RESOURCE_GROUP} \
--template-file "<installation_directory>/02_storage.json" \
--parameters vhdBlobURL="${VHD_BLOB_URL}" \ 1
```

```
--parameters baseName="${INFRA_ID}" \ 2
--parameters storageAccount="${CLUSTER_NAME}sa" \ 3
--parameters architecture="<architecture>" 4
```

- 1 The blob URL of the RHCOS VHD to be used to create master and worker machines.
- 2 The base name to be used in resource names; this is usually the cluster's infrastructure ID.
- 3 The name of your Azure storage account.
- 4 Specify the system architecture. Valid values are **x64** (default) or **Arm64**.

7.11.13.1. ARM template for image storage

You can use the following Azure Resource Manager (ARM) template to deploy the stored Red Hat Enterprise Linux CoreOS (RHCOS) image that you need for your OpenShift Container Platform cluster:

Example 7.86. 02_storage.json ARM template

```
{
  "$schema": "https://schema.management.azure.com/schemas/2019-04-01/deploymentTemplate.json#",
  "contentVersion": "1.0.0.0",
  "parameters": {
    "architecture": {
      "type": "string",
      "metadata": {
        "description": "The architecture of the Virtual Machines"
      },
      "defaultValue": "x64",
      "allowedValues": [
        "Arm64",
        "x64"
      ]
    },
    "baseName": {
      "type": "string",
      "minLength": 1,
      "metadata": {
        "description": "Base name to be used in resource names (usually the cluster's Infra ID)"
      }
    },
    "storageAccount": {
      "type": "string",
      "metadata": {
        "description": "The Storage Account name"
      }
    },
    "vhdBlobURL": {
      "type": "string",
      "metadata": {
        "description": "URL pointing to the blob where the VHD to be used to create master and worker machines is located"
      }
    }
  }
}
```

```


    }
  },
  "variables": {
    "location": "[resourceGroup().location]",
    "galleryName": "[concat('gallery_', replace(parameters('baseName'), '-', '_'))]",
    "imageName": "[parameters('baseName')]",
    "imageNameGen2": "[concat(parameters('baseName'), '-gen2')]",
    "imageRelease": "1.0.0"
  },
  "resources": [
    {
      "apiVersion": "2021-10-01",
      "type": "Microsoft.Compute/galleries",
      "name": "[variables('galleryName')]",
      "location": "[variables('location')]",
      "resources": [
        {
          "apiVersion": "2021-10-01",
          "type": "images",
          "name": "[variables('imageName')]",
          "location": "[variables('location')]",
          "dependsOn": [
            "[variables('galleryName')]"
          ],
          "properties": {
            "architecture": "[parameters('architecture')]",
            "hyperVGeneration": "V1",
            "identifier": {
              "offer": "rhcos",
              "publisher": "RedHat",
              "sku": "basic"
            },
            "osState": "Generalized",
            "osType": "Linux"
          },
          "resources": [
            {
              "apiVersion": "2021-10-01",
              "type": "versions",
              "name": "[variables('imageRelease')]",
              "location": "[variables('location')]",
              "dependsOn": [
                "[variables('imageName')]"
              ],
              "properties": {
                "publishingProfile": {
                  "storageAccountType": "Standard_LRS",
                  "targetRegions": [
                    {
                      "name": "[variables('location')]",
                      "regionalReplicaCount": "1"
                    }
                  ]
                }
              },
              "storageProfile": {
                "osDiskImage": {

```

```

        "source": {
          "id": "[resourceId('Microsoft.Storage/storageAccounts',
parameters('storageAccount'))]",
          "uri": "[parameters('vhdBlobURL')]"
        }
      }
    }
  }
}
],
},
{
  "apiVersion": "2021-10-01",
  "type": "images",
  "name": "[variables('imageNameGen2')]",
  "location": "[variables('location')]",
  "dependsOn": [
    "[variables('galleryName')]"
  ],
  "properties": {
    "architecture": "[parameters('architecture')]",
    "hyperVGeneration": "V2",
    "identifier": {
      "offer": "rhcos-gen2",
      "publisher": "RedHat-gen2",
      "sku": "gen2"
    },
    "osState": "Generalized",
    "osType": "Linux"
  },
  "resources": [
    {
      "apiVersion": "2021-10-01",
      "type": "versions",
      "name": "[variables('imageRelease')]",
      "location": "[variables('location')]",
      "dependsOn": [
        "[variables('imageNameGen2')]"
      ],
      "properties": {
        "publishingProfile": {
          "storageAccountType": "Standard_LRS",
          "targetRegions": [
            {
              "name": "[variables('location')]",
              "regionalReplicaCount": "1"
            }
          ]
        }
      },
      "storageProfile": {
        "osDiskImage": {
          "source": {
            "id": "[resourceId('Microsoft.Storage/storageAccounts',
parameters('storageAccount'))]",
            "uri": "[parameters('vhdBlobURL')]"
          }
        }
      }
    }
  ]
}

```




7.11.14. Networking requirements for user-provisioned infrastructure

All the Red Hat Enterprise Linux CoreOS (RHCOS) machines require networking to be configured in **initramfs** during boot to fetch their Ignition config files.

7.11.14.1. Network connectivity requirements

You must configure the network connectivity between machines to allow OpenShift Container Platform cluster components to communicate. Each machine must be able to resolve the hostnames of all other machines in the cluster.

This section provides details about the ports that are required.



IMPORTANT

In connected OpenShift Container Platform environments, all nodes are required to have internet access to pull images for platform containers and provide telemetry data to Red Hat.

Table 7.27. Ports used for all-machine to all-machine communications

Protocol	Port	Description
ICMP	N/A	Network reachability tests
TCP	1936	Metrics
	9000-9999	Host level services, including the node exporter on ports 9100-9101 and the Cluster Version Operator on port 9099 .
	10250-10259	The default ports that Kubernetes reserves
UDP	4789	VXLAN
	6081	Geneve
	9000-9999	Host level services, including the node exporter on ports 9100-9101 .

Protocol	Port	Description
	500	IPsec IKE packets
	4500	IPsec NAT-T packets
	123	Network Time Protocol (NTP) on UDP port 123 If an external NTP time server is configured, you must open UDP port 123 .
TCP/UDP	30000-32767	Kubernetes node port
ESP	N/A	IPsec Encapsulating Security Payload (ESP)

Table 7.28. Ports used for all-machine to control plane communications


Protocol	Port	Description
TCP	6443	Kubernetes API

Table 7.29. Ports used for control plane machine to control plane machine communications

Protocol	Port	Description
TCP	2379-2380	etcd server and peer ports

7.11.15. Creating networking and load balancing components in Azure

You must configure networking and load balancing in Microsoft Azure for your OpenShift Container Platform cluster to use. One way to create these components is to modify the provided Azure Resource Manager (ARM) template.



NOTE

If you do not use the provided ARM template to create your Azure infrastructure, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- Configure an Azure account.
- Generate the Ignition config files for your cluster.
- Create and configure a VNet and associated subnets in Azure.

Procedure

1. Copy the template from the **ARM template for the network and load balancers** section of this topic and save it as **03_infra.json** in your cluster's installation directory. This template describes the networking and load balancing objects that your cluster requires.
2. Create the deployment by using the **az** CLI:

```
$ az deployment group create -g ${RESOURCE_GROUP} \
  --template-file "<installation_directory>/03_infra.json" \
  --parameters privateDNSZoneName="${CLUSTER_NAME}.${BASE_DOMAIN}" \ 1
  --parameters baseName="${INFRA_ID}" 2
```

1 The name of the private DNS zone.

2 The base name to be used in resource names; this is usually the cluster's infrastructure ID.

3. Create an **api** DNS record in the public zone for the API public load balancer. The **\${BASE_DOMAIN_RESOURCE_GROUP}** variable must point to the resource group where the public DNS zone exists.

- a. Export the following variable:

```
$ export PUBLIC_IP=`az network public-ip list -g ${RESOURCE_GROUP} --query "[?
name=='${INFRA_ID}-master-pip'] | [0].ipAddress" -o tsv`
```

- b. Create the **api** DNS record in a new public zone:

```
$ az network dns record-set a add-record -g ${BASE_DOMAIN_RESOURCE_GROUP} -
z ${CLUSTER_NAME}.${BASE_DOMAIN} -n api -a ${PUBLIC_IP} --ttl 60
```

If you are adding the cluster to an existing public zone, you can create the **api** DNS record in it instead:

```
$ az network dns record-set a add-record -g ${BASE_DOMAIN_RESOURCE_GROUP} -
z ${BASE_DOMAIN} -n api.${CLUSTER_NAME} -a ${PUBLIC_IP} --ttl 60
```

7.11.15.1. ARM template for the network and load balancers

You can use the following Azure Resource Manager (ARM) template to deploy the networking objects and load balancers that you need for your OpenShift Container Platform cluster:

Example 7.87. 03_infra.json ARM template

```
{
  "$schema" : "https://schema.management.azure.com/schemas/2015-01-
01/deploymentTemplate.json#",
  "contentVersion" : "1.0.0.0",
  "parameters" : {
    "baseName" : {
      "type" : "string",
      "minLength" : 1,
      "metadata" : {
        "description" : "Base name to be used in resource names (usually the cluster's Infra ID)"
      }
    }
  }
}
```

```

},
"vnetBaseName": {
  "type": "string",
  "defaultValue": "",
  "metadata": {
    "description": "The specific customer vnet's base name (optional)"
  }
},
"privateDNSZoneName": {
  "type": "string",
  "metadata": {
    "description": "Name of the private DNS zone"
  }
}
},
"variables": {
  "location": "[resourceGroup().location]",
  "virtualNetworkName": "[concat(if(not(empty(parameters('vnetBaseName'))),
parameters('vnetBaseName'), parameters('baseName')), '-vnet')]",
  "virtualNetworkID": "[resourceId('Microsoft.Network/virtualNetworks',
variables('virtualNetworkName'))]",
  "masterSubnetName": "[concat(if(not(empty(parameters('vnetBaseName'))),
parameters('vnetBaseName'), parameters('baseName')), '-master-subnet')]",
  "masterSubnetRef": "[concat(variables('virtualNetworkID'), '/subnets/',
variables('masterSubnetName'))]",
  "masterPublicIpAddressName": "[concat(parameters('baseName'), '-master-pip')]",
  "masterPublicIpAddressID": "[resourceId('Microsoft.Network/publicIPAddresses',
variables('masterPublicIpAddressName'))]",
  "masterLoadBalancerName": "[parameters('baseName')]",
  "masterLoadBalancerID": "[resourceId('Microsoft.Network/loadBalancers',
variables('masterLoadBalancerName'))]",
  "internalLoadBalancerName": "[concat(parameters('baseName'), '-internal-lb')]",
  "internalLoadBalancerID": "[resourceId('Microsoft.Network/loadBalancers',
variables('internalLoadBalancerName'))]",
  "skuName": "Standard"
},
"resources": [
{
  "apiVersion": "2018-12-01",
  "type": "Microsoft.Network/publicIPAddresses",
  "name": "[variables('masterPublicIpAddressName')]",
  "location": "[variables('location')]",
  "sku": {
    "name": "[variables('skuName')]"
  },
  "properties": {
    "publicIPAllocationMethod": "Static",
    "dnsSettings": {
      "domainNameLabel": "[variables('masterPublicIpAddressName')]"
    }
  }
},
{
  "apiVersion": "2018-12-01",
  "type": "Microsoft.Network/loadBalancers",
  "name": "[variables('masterLoadBalancerName')]",

```

```

"location" : "[variables('location')]",
"sku": {
  "name": "[variables('skuName')]"
},
"dependsOn" : [
  "[concat('Microsoft.Network/publicIPAddresses/', variables('masterPublicIpAddressName'))]"
],
"properties" : {
  "frontendIPConfigurations" : [
    {
      "name" : "public-lb-ip-v4",
      "properties" : {
        "publicIPAddress" : {
          "id" : "[variables('masterPublicIpAddressID')]"
        }
      }
    }
  ],
  "backendAddressPools" : [
    {
      "name" : "[variables('masterLoadBalancerName')]"
    }
  ],
  "loadBalancingRules" : [
    {
      "name" : "api-internal",
      "properties" : {
        "frontendIPConfiguration" : {
          "id" : "[concat(variables('masterLoadBalancerID'), '/frontendIPConfigurations/public-lb-ip-
v4')]"
        },
        "backendAddressPool" : {
          "id" : "[concat(variables('masterLoadBalancerID'), '/backendAddressPools/',
variables('masterLoadBalancerName'))]"
        },
        "protocol" : "Tcp",
        "loadDistribution" : "Default",
        "idleTimeoutInMinutes" : 30,
        "frontendPort" : 6443,
        "backendPort" : 6443,
        "probe" : {
          "id" : "[concat(variables('masterLoadBalancerID'), '/probes/api-internal-probe')]"
        }
      }
    }
  ],
  "probes" : [
    {
      "name" : "api-internal-probe",
      "properties" : {
        "protocol" : "Https",
        "port" : 6443,
        "requestPath" : "/readyz",
        "intervalInSeconds" : 10,
        "numberOfProbes" : 3
      }
    }
  ]
}

```

```

    }
  ]
}
},
{
  "apiVersion" : "2018-12-01",
  "type" : "Microsoft.Network/loadBalancers",
  "name" : "[variables('internalLoadBalancerName')]",
  "location" : "[variables('location')]",
  "sku": {
    "name": "[variables('skuName')]"
  },
  "properties" : {
    "frontendIPConfigurations" : [
      {
        "name" : "internal-lb-ip",
        "properties" : {
          "privateIPAllocationMethod" : "Dynamic",
          "subnet" : {
            "id" : "[variables('masterSubnetRef')]"
          },
          "privateIPAddressVersion" : "IPv4"
        }
      }
    ],
    "backendAddressPools" : [
      {
        "name" : "internal-lb-backend"
      }
    ],
    "loadBalancingRules" : [
      {
        "name" : "api-internal",
        "properties" : {
          "frontendIPConfiguration" : {
            "id" : "[concat(variables('internalLoadBalancerID'), '/frontendIPConfigurations/internal-lb-
ip')]"
          },
          "frontendPort" : 6443,
          "backendPort" : 6443,
          "enableFloatingIP" : false,
          "idleTimeoutInMinutes" : 30,
          "protocol" : "Tcp",
          "enableTcpReset" : false,
          "loadDistribution" : "Default",
          "backendAddressPool" : {
            "id" : "[concat(variables('internalLoadBalancerID'), '/backendAddressPools/internal-lb-
backend')]"
          },
          "probe" : {
            "id" : "[concat(variables('internalLoadBalancerID'), '/probes/api-internal-probe')]"
          }
        }
      }
    ],
    {
      "name" : "sint",

```

```

    "properties" : {
      "frontendIPConfiguration" : {
        "id" : "[concat(variables('internalLoadBalancerID'), '/frontendIPConfigurations/internal-lb-
ip)']"
      },
      "frontendPort" : 22623,
      "backendPort" : 22623,
      "enableFloatingIP" : false,
      "idleTimeoutInMinutes" : 30,
      "protocol" : "Tcp",
      "enableTcpReset" : false,
      "loadDistribution" : "Default",
      "backendAddressPool" : {
        "id" : "[concat(variables('internalLoadBalancerID'), '/backendAddressPools/internal-lb-
backend)']"
      },
      "probe" : {
        "id" : "[concat(variables('internalLoadBalancerID'), '/probes/sint-probe)']"
      }
    }
  },
  "probes" : [
    {
      "name" : "api-internal-probe",
      "properties" : {
        "protocol" : "Https",
        "port" : 6443,
        "requestPath" : "/readyz",
        "intervalInSeconds" : 10,
        "numberOfProbes" : 3
      }
    },
    {
      "name" : "sint-probe",
      "properties" : {
        "protocol" : "Https",
        "port" : 22623,
        "requestPath" : "/healthz",
        "intervalInSeconds" : 10,
        "numberOfProbes" : 3
      }
    }
  ]
},
{
  "apiVersion": "2018-09-01",
  "type": "Microsoft.Network/privateDnsZones/A",
  "name": "[concat(parameters('privateDNSZoneName'), '/api')]",
  "location": "[variables('location')]",
  "dependsOn" : [
    "[concat('Microsoft.Network/loadBalancers/', variables('internalLoadBalancerName'))]"
  ],
  "properties": {
    "ttl": 60,

```

```

    "aRecords": [
      {
        "ipv4Address": "[reference(variables('internalLoadBalancerName')).frontendIPConfigurations[0].properties.privateIP
Address]"
      }
    ]
  },
  {
    "apiVersion": "2018-09-01",
    "type": "Microsoft.Network/privateDnsZones/A",
    "name": "[concat(parameters('privateDNSZoneName'), '/api-int')]",
    "location": "[variables('location')]",
    "dependsOn": [
      "[concat('Microsoft.Network/loadBalancers/', variables('internalLoadBalancerName'))]"
    ],
    "properties": {
      "ttl": 60,
      "aRecords": [
        {
          "ipv4Address": "[reference(variables('internalLoadBalancerName')).frontendIPConfigurations[0].properties.privateIP
Address]"
        }
      ]
    }
  }
]
}

```

7.11.16. Creating the bootstrap machine in Azure

You must create the bootstrap machine in Microsoft Azure to use during OpenShift Container Platform cluster initialization. One way to create this machine is to modify the provided Azure Resource Manager (ARM) template.



NOTE

If you do not use the provided ARM template to create your bootstrap machine, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- Configure an Azure account.
- Generate the Ignition config files for your cluster.
- Create and configure a VNet and associated subnets in Azure.
- Create and configure networking and load balancers in Azure.

- Create control plane and compute roles.

Procedure

1. Copy the template from the **ARM template for the bootstrap machine** section of this topic and save it as **04_bootstrap.json** in your cluster's installation directory. This template describes the bootstrap machine that your cluster requires.
2. Export the bootstrap URL variable:

```
$ bootstrap_url_expiry=`date -u -d "10 hours" '+%Y-%m-%dT%H:%MZ'`
```

```
$ export BOOTSTRAP_URL=`az storage blob generate-sas -c 'files' -n 'bootstrap.ign' --https-only --full-uri --permissions r --expiry $bootstrap_url_expiry --account-name ${CLUSTER_NAME}sa --account-key ${ACCOUNT_KEY} -o tsv`
```

3. Export the bootstrap ignition variable:

```
$ export BOOTSTRAP_IGNITION=`jq -rcnM --arg v "3.2.0" --arg url ${BOOTSTRAP_URL} '{ignition:{version:$v,config:{replace:{source:$url}}}}' | base64 | tr -d '\n'`
```

4. Create the deployment by using the **az** CLI:

```
$ az deployment group create -g ${RESOURCE_GROUP} \
  --template-file "<installation_directory>/04_bootstrap.json" \
  --parameters bootstrapIgnition="${BOOTSTRAP_IGNITION}" \ 1
  --parameters baseName="${INFRA_ID}" \ 2
  --parameter bootstrapVMSize="Standard_D4s_v3" 3
```

- 1** The bootstrap Ignition content for the bootstrap cluster.
- 2** The base name to be used in resource names; this is usually the cluster's infrastructure ID.
- 3** Optional: Specify the size of the bootstrap VM. Use a VM size compatible with your specified architecture. If this value is not defined, the default value from the template is set.

7.11.16.1. ARM template for the bootstrap machine

You can use the following Azure Resource Manager (ARM) template to deploy the bootstrap machine that you need for your OpenShift Container Platform cluster:

Example 7.88. 04_bootstrap.json ARM template

```
{
  "$schema" : "https://schema.management.azure.com/schemas/2015-01-01/deploymentTemplate.json#",
  "contentVersion" : "1.0.0.0",
  "parameters" : {
    "baseName" : {
      "type" : "string",
      "minLength" : 1,
```



```

"metadata" : {
  "description" : "Base name to be used in resource names (usually the cluster's Infra ID)"
}
},
"vnetBaseName": {
  "type": "string",
  "defaultValue": "",
  "metadata" : {
    "description" : "The specific customer vnet's base name (optional)"
  }
},
"bootstrapIgnition" : {
  "type" : "string",
  "minLength" : 1,
  "metadata" : {
    "description" : "Bootstrap ignition content for the bootstrap cluster"
  }
},
"sshKeyData" : {
  "type" : "securestring",
  "defaultValue" : "Unused",
  "metadata" : {
    "description" : "Unused"
  }
},
"bootstrapVMSize" : {
  "type" : "string",
  "defaultValue" : "Standard_D4s_v3",
  "metadata" : {
    "description" : "The size of the Bootstrap Virtual Machine"
  }
},
"hyperVGen": {
  "type": "string",
  "metadata": {
    "description": "VM generation image to use"
  },
  "defaultValue": "V2",
  "allowedValues": [
    "V1",
    "V2"
  ]
},
"variables" : {
  "location" : "[resourceGroup().location]",
  "virtualNetworkName" : "[concat(if(not(empty(parameters('vnetBaseName'))),
parameters('vnetBaseName'), parameters('baseName')), '-vnet')]",
  "virtualNetworkID" : "[resourceId('Microsoft.Network/virtualNetworks',
variables('virtualNetworkName'))]",
  "masterSubnetName" : "[concat(if(not(empty(parameters('vnetBaseName'))),
parameters('vnetBaseName'), parameters('baseName')), '-master-subnet')]",
  "masterSubnetRef" : "[concat(variables('virtualNetworkID'), '/subnets/',
variables('masterSubnetName'))]",
  "masterLoadBalancerName" : "[parameters('baseName')]",
  "internalLoadBalancerName" : "[concat(parameters('baseName'), '-internal-lb')]",

```

```

"sshKeyPath" : "/home/core/.ssh/authorized_keys",
"identityName" : "[concat(parameters('baseName'), '-identity')]",
"vmName" : "[concat(parameters('baseName'), '-bootstrap')]",
"nicName" : "[concat(variables('vmName'), '-nic')]",
"galleryName": "[concat('gallery_', replace(parameters('baseName'), '-', '_'))]",
"imageName" : "[concat(parameters('baseName'), if(equals(parameters('hyperVGen'), 'V2'), '-gen2', ''))]",
"clusterNsgName" : "[concat(if(not(empty(parameters('vnetBaseName'))), parameters('vnetBaseName'), parameters('baseName')), '-nsg')]",
"sshPublicIpAddressName" : "[concat(variables('vmName'), '-ssh-pip')]"
},
"resources" : [
{
"apiVersion" : "2018-12-01",
"type" : "Microsoft.Network/publicIPAddresses",
"name" : "[variables('sshPublicIpAddressName')]",
"location" : "[variables('location')]",
"sku": {
"name": "Standard"
},
"properties" : {
"publicIPAllocationMethod" : "Static",
"dnsSettings" : {
"domainNameLabel" : "[variables('sshPublicIpAddressName')]"
}
}
},
{
"apiVersion" : "2018-06-01",
"type" : "Microsoft.Network/networkInterfaces",
"name" : "[variables('nicName')]",
"location" : "[variables('location')]",
"dependsOn" : [
"[resourceId('Microsoft.Network/publicIPAddresses', variables('sshPublicIpAddressName'))]"
],
"properties" : {
"ipConfigurations" : [
{
"name" : "pipConfig",
"properties" : {
"privateIPAllocationMethod" : "Dynamic",
"publicIPAddress": {
"id": "[resourceId('Microsoft.Network/publicIPAddresses', variables('sshPublicIpAddressName'))]"
},
"subnet" : {
"id" : "[variables('masterSubnetRef')]"
},
"loadBalancerBackendAddressPools" : [
{
"id" : "[concat('/subscriptions/', subscription().subscriptionId, '/resourceGroups/', resourceGroup().name, '/providers/Microsoft.Network/loadBalancers/', variables('masterLoadBalancerName'), '/backendAddressPools/', variables('masterLoadBalancerName'))]"
},
{

```

```

        "id" : "[concat('/subscriptions/', subscription().subscriptionId, '/resourceGroups/',
resourceGroup().name, '/providers/Microsoft.Network/loadBalancers/',
variables('internalLoadBalancerName'), '/backendAddressPools/internal-lb-backend')]"
    }
  ]
}
}
]
},
{
  "apiVersion" : "2018-06-01",
  "type" : "Microsoft.Compute/virtualMachines",
  "name" : "[variables('vmName')]",
  "location" : "[variables('location')]",
  "identity" : {
    "type" : "userAssigned",
    "userAssignedIdentities" : {
      "[resourceID('Microsoft.ManagedIdentity/userAssignedIdentities/',
variables('identityName'))]" : {}
    }
  },
  "dependsOn" : [
    "[concat('Microsoft.Network/networkInterfaces/', variables('nicName'))]"
  ],
  "properties" : {
    "hardwareProfile" : {
      "vmSize" : "[parameters('bootstrapVMSize')]"
    },
    "osProfile" : {
      "computerName" : "[variables('vmName')]",
      "adminUsername" : "core",
      "adminPassword" : "NotActuallyApplied!",
      "customData" : "[parameters('bootstrapIgnition')]",
      "linuxConfiguration" : {
        "disablePasswordAuthentication" : false
      }
    },
    "storageProfile" : {
      "imageReference" : {
        "id" : "[resourceId('Microsoft.Compute/galleries/images', variables('galleryName'),
variables('imageName'))]"
      },
      "osDisk" : {
        "name" : "[concat(variables('vmName'), '_OSDisk')]",
        "osType" : "Linux",
        "createOption" : "FromImage",
        "managedDisk" : {
          "storageAccountType" : "Premium_LRS"
        },
        "diskSizeGB" : 100
      }
    },
    "networkProfile" : {
      "networkInterfaces" : [
        {

```

```

        "id" : "[resourceId('Microsoft.Network/networkInterfaces', variables('nicName'))]"
      }
    ]
  }
},
{
  "apiVersion" : "2018-06-01",
  "type": "Microsoft.Network/networkSecurityGroups/securityRules",
  "name" : "[concat(variables('clusterNsgName'), '/bootstrap_ssh_in')]",
  "location" : "[variables('location')]",
  "dependsOn" : [
    "[resourceId('Microsoft.Compute/virtualMachines', variables('vmName'))]"
  ],
  "properties": {
    "protocol" : "Tcp",
    "sourcePortRange" : "*",
    "destinationPortRange" : "22",
    "sourceAddressPrefix" : "*",
    "destinationAddressPrefix" : "*",
    "access" : "Allow",
    "priority" : 100,
    "direction" : "Inbound"
  }
}
]
}

```

7.11.17. Creating the control plane machines in Azure

You must create the control plane machines in Microsoft Azure for your cluster to use. One way to create these machines is to modify the provided Azure Resource Manager (ARM) template.



NOTE

By default, Microsoft Azure places control plane machines and compute machines in a pre-set availability zone. You can manually set an availability zone for a compute node or control plane node. To do this, modify a vendor's Azure Resource Manager (ARM) template by specifying each of your availability zones in the **zones** parameter of the virtual machine resource.

If you do not use the provided ARM template to create your control plane machines, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, consider contacting Red Hat support with your installation logs.

Prerequisites

- Configure an Azure account.
- Generate the Ignition config files for your cluster.
- Create and configure a VNet and associated subnets in Azure.

- Create and configure networking and load balancers in Azure.
- Create control plane and compute roles.
- Create the bootstrap machine.

Procedure

1. Copy the template from the **ARM template for control plane machines** section of this topic and save it as **05_masters.json** in your cluster's installation directory. This template describes the control plane machines that your cluster requires.
2. Export the following variable needed by the control plane machine deployment:

```
$ export MASTER_IGNITION=`cat <installation_directory>/master.ign | base64 | tr -d "\n`
```

3. Create the deployment by using the **az** CLI:

```
$ az deployment group create -g ${RESOURCE_GROUP} \
  --template-file "<installation_directory>/05_masters.json" \
  --parameters masterIgnition="${MASTER_IGNITION}" \ 1
  --parameters baseName="${INFRA_ID}" \ 2
  --parameters masterVMSize="Standard_D8s_v3" 3
```

- 1** The Ignition content for the control plane nodes.
- 2** The base name to be used in resource names; this is usually the cluster's infrastructure ID.
- 3** Optional: Specify the size of the Control Plane VM. Use a VM size compatible with your specified architecture. If this value is not defined, the default value from the template is set.

7.11.17.1. ARM template for control plane machines

You can use the following Azure Resource Manager (ARM) template to deploy the control plane machines that you need for your OpenShift Container Platform cluster:

Example 7.89. 05_masters.json ARM template

```
{
  "$schema" : "https://schema.management.azure.com/schemas/2015-01-01/deploymentTemplate.json#",
  "contentVersion" : "1.0.0.0",
  "parameters" : {
    "baseName" : {
      "type" : "string",
      "minLength" : 1,
      "metadata" : {
        "description" : "Base name to be used in resource names (usually the cluster's Infra ID)"
      }
    },
    "vnetBaseName": {
      "type": "string",
```

```
"defaultValue": "",
"metadata" : {
  "description" : "The specific customer vnet's base name (optional)"
}
},
"masterIgnition" : {
  "type" : "string",
  "metadata" : {
    "description" : "Ignition content for the master nodes"
  }
},
"numberOfMasters" : {
  "type" : "int",
  "defaultValue" : 3,
  "minValue" : 2,
  "maxValue" : 30,
  "metadata" : {
    "description" : "Number of OpenShift masters to deploy"
  }
},
"sshKeyData" : {
  "type" : "securestring",
  "defaultValue" : "Unused",
  "metadata" : {
    "description" : "Unused"
  }
},
"privateDNSZoneName" : {
  "type" : "string",
  "defaultValue" : "",
  "metadata" : {
    "description" : "unused"
  }
},
"masterVMSize" : {
  "type" : "string",
  "defaultValue" : "Standard_D8s_v3",
  "metadata" : {
    "description" : "The size of the Master Virtual Machines"
  }
},
"diskSizeGB" : {
  "type" : "int",
  "defaultValue" : 1024,
  "metadata" : {
    "description" : "Size of the Master VM OS disk, in GB"
  }
},
"hyperVGen": {
  "type": "string",
  "metadata": {
    "description": "VM generation image to use"
  },
  "defaultValue": "V2",
  "allowedValues": [
    "V1",
```

```

    "V2"
  ]
}
},
"variables" : {
  "location" : "[resourceGroup().location]",
  "virtualNetworkName" : "[concat(if(not(empty(parameters('vnetBaseName'))),
parameters('vnetBaseName'), parameters('baseName')), '-vnet')]",
  "virtualNetworkID" : "[resourceId('Microsoft.Network/virtualNetworks',
variables('virtualNetworkName'))]",
  "masterSubnetName" : "[concat(if(not(empty(parameters('vnetBaseName'))),
parameters('vnetBaseName'), parameters('baseName')), '-master-subnet')]",
  "masterSubnetRef" : "[concat(variables('virtualNetworkID'), '/subnets/',
variables('masterSubnetName'))]",
  "masterLoadBalancerName" : "[parameters('baseName')]",
  "internalLoadBalancerName" : "[concat(parameters('baseName'), '-internal-lb')]",
  "sshKeyPath" : "/home/core/.ssh/authorized_keys",
  "identityName" : "[concat(parameters('baseName'), '-identity')]",
  "galleryName": "[concat('gallery_', replace(parameters('baseName'), '-', '_'))]",
  "imageName" : "[concat(parameters('baseName'), if(equals(parameters('hyperVGen'), 'V2'), '-
gen2', ''))]",
  "copy" : [
    {
      "name" : "vmNames",
      "count" : "[parameters('numberOfMasters')]",
      "input" : "[concat(parameters('baseName'), '-master-', copyIndex('vmNames'))]"
    }
  ]
},
"resources" : [
  {
    "apiVersion" : "2018-06-01",
    "type" : "Microsoft.Network/networkInterfaces",
    "copy" : {
      "name" : "nicCopy",
      "count" : "[length(variables('vmNames'))]"
    },
    "name" : "[concat(variables('vmNames')[copyIndex()], '-nic')]",
    "location" : "[variables('location')]",
    "properties" : {
      "ipConfigurations" : [
        {
          "name" : "pipConfig",
          "properties" : {
            "privateIPAllocationMethod" : "Dynamic",
            "subnet" : {
              "id" : "[variables('masterSubnetRef')]"
            },
            "loadBalancerBackendAddressPools" : [
              {
                "id" : "[concat('/subscriptions/', subscription().subscriptionId, '/resourceGroups/',
resourceGroup().name, '/providers/Microsoft.Network/loadBalancers/',
variables('masterLoadBalancerName'), '/backendAddressPools/',
variables('masterLoadBalancerName'))]"
              },
            ]
          }
        }
      ]
    }
  }
]
}

```

```
        "id" : "[concat('/subscriptions/', subscription().subscriptionId, '/resourceGroups/',
resourceGroup().name, '/providers/Microsoft.Network/loadBalancers/',
variables('internalLoadBalancerName'), '/backendAddressPools/internal-lb-backend')]"
    }
  ]
}
}
],
},
{
  "apiVersion" : "2018-06-01",
  "type" : "Microsoft.Compute/virtualMachines",
  "copy" : {
    "name" : "vmCopy",
    "count" : "[length(variables('vmNames'))]"
  },
  "name" : "[variables('vmNames')[copyIndex()]]",
  "location" : "[variables('location')]",
  "identity" : {
    "type" : "userAssigned",
    "userAssignedIdentities" : {
      "[resourceID('Microsoft.ManagedIdentity/userAssignedIdentities/',
variables('identityName'))]" : {}
    }
  },
  "dependsOn" : [
    "[concat('Microsoft.Network/networkInterfaces/', concat(variables('vmNames')[copyIndex()], '-
nic'))]"
  ],
  "properties" : {
    "hardwareProfile" : {
      "vmSize" : "[parameters('masterVMSize')]"
    },
    "osProfile" : {
      "computerName" : "[variables('vmNames')[copyIndex()]]",
      "adminUsername" : "core",
      "adminPassword" : "NotActuallyApplied!",
      "customData" : "[parameters('masterIgnition')]",
      "linuxConfiguration" : {
        "disablePasswordAuthentication" : false
      }
    },
    "storageProfile" : {
      "imageReference": {
        "id": "[resourceId('Microsoft.Compute/galleries/images', variables('galleryName'),
variables('imageName'))]"
      },
      "osDisk" : {
        "name": "[concat(variables('vmNames')[copyIndex()], '_OSDisk')]",
        "osType" : "Linux",
        "createOption" : "FromImage",
        "caching" : "ReadOnly",
        "writeAcceleratorEnabled" : false,
        "managedDisk": {
          "storageAccountType": "Premium_LRS"
        }
      }
    }
  }
}
```



```

    },
    "diskSizeGB" : "[parameters('diskSizeGB')]"
  }
},
"networkProfile" : {
  "networkInterfaces" : [
    {
      "id" : "[resourceId('Microsoft.Network/networkInterfaces', concat(variables('vmNames')
[copyIndex()], '-nic'))]",
      "properties" : {
        "primary" : false
      }
    }
  ]
}
}
]
}
}
]
}

```

7.11.18. Wait for bootstrap completion and remove bootstrap resources in Azure

After you create all of the required infrastructure in Microsoft Azure, wait for the bootstrap process to complete on the machines that you provisioned by using the Ignition config files that you generated with the installation program.

Prerequisites

- Configure an Azure account.
- Generate the Ignition config files for your cluster.
- Create and configure a VNet and associated subnets in Azure.
- Create and configure networking and load balancers in Azure.
- Create control plane and compute roles.
- Create the bootstrap machine.
- Create the control plane machines.

Procedure

1. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install wait-for bootstrap-complete --dir <installation_directory> \ 1
--log-level info 2
```

- 1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.
- 2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

If the command exits without a **FATAL** warning, your production control plane has initialized.

2. Delete the bootstrap resources:

```
$ az network nsg rule delete -g ${RESOURCE_GROUP} --nsg-name ${INFRA_ID}-nsg --name bootstrap_ssh_in
$ az vm stop -g ${RESOURCE_GROUP} --name ${INFRA_ID}-bootstrap
$ az vm deallocate -g ${RESOURCE_GROUP} --name ${INFRA_ID}-bootstrap
$ az vm delete -g ${RESOURCE_GROUP} --name ${INFRA_ID}-bootstrap --yes
$ az disk delete -g ${RESOURCE_GROUP} --name ${INFRA_ID}-bootstrap OSDisk --no-wait --yes
$ az network nic delete -g ${RESOURCE_GROUP} --name ${INFRA_ID}-bootstrap-nic --no-wait
$ az storage blob delete --account-key ${ACCOUNT_KEY} --account-name ${CLUSTER_NAME}sa --container-name files --name bootstrap.ign
$ az network public-ip delete -g ${RESOURCE_GROUP} --name ${INFRA_ID}-bootstrap-ssh-pip
```




NOTE

If you do not delete the bootstrap server, installation may not succeed due to API traffic being routed to the bootstrap server.

7.11.19. Creating additional worker machines in Azure

You can create worker machines in Microsoft Azure for your cluster to use by launching individual instances discretely or by automated processes outside the cluster, such as auto scaling groups. You can also take advantage of the built-in cluster scaling mechanisms and the machine API in OpenShift Container Platform.



NOTE

If you are installing a three-node cluster, skip this step. A three-node cluster consists of three control plane machines, which also act as compute machines.

In this example, you manually launch one instance by using the Azure Resource Manager (ARM) template. Additional instances can be launched by including additional resources of type **06_workers.json** in the file.



NOTE

By default, Microsoft Azure places control plane machines and compute machines in a pre-set availability zone. You can manually set an availability zone for a compute node or control plane node. To do this, modify a vendor's ARM template by specifying each of your availability zones in the **zones** parameter of the virtual machine resource.

If you do not use the provided ARM template to create your control plane machines, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, consider contacting Red Hat support with your installation logs.

Prerequisites

- Configure an Azure account.

- Generate the Ignition config files for your cluster.
- Create and configure a VNet and associated subnets in Azure.
- Create and configure networking and load balancers in Azure.
- Create control plane and compute roles.
- Create the bootstrap machine.
- Create the control plane machines.

Procedure

1. Copy the template from the **ARM template for worker machines** section of this topic and save it as **06_workers.json** in your cluster's installation directory. This template describes the worker machines that your cluster requires.
2. Export the following variable needed by the worker machine deployment:

```
$ export WORKER_IGNITION=`cat <installation_directory>/worker.ign | base64 | tr -d "\n`
```

3. Create the deployment by using the **az** CLI:

```
$ az deployment group create -g ${RESOURCE_GROUP} \
  --template-file "<installation_directory>/06_workers.json" \
  --parameters workerIgnition="${WORKER_IGNITION}" \ 1
  --parameters baseName="${INFRA_ID}" \ 2
  --parameters nodeVMSize="Standard_D4s_v3" 3
```

- 1** The Ignition content for the worker nodes.
- 2** The base name to be used in resource names; this is usually the cluster's infrastructure ID.
- 3** Optional: Specify the size of the compute node VM. Use a VM size compatible with your specified architecture. If this value is not defined, the default value from the template is set.

7.11.19.1. ARM template for worker machines

You can use the following Azure Resource Manager (ARM) template to deploy the worker machines that you need for your OpenShift Container Platform cluster:

Example 7.90. 06_workers.json ARM template

```
{
  "$schema" : "https://schema.management.azure.com/schemas/2015-01-01/deploymentTemplate.json#",
  "contentVersion" : "1.0.0.0",
  "parameters" : {
    "baseName" : {
      "type" : "string",
      "minLength" : 1,
      "metadata" : {
```

```

    "description" : "Base name to be used in resource names (usually the cluster's Infra ID)"
  }
},
"vnetBaseName": {
  "type": "string",
  "defaultValue": "",
  "metadata" : {
    "description" : "The specific customer vnet's base name (optional)"
  }
},
"workerIgnition" : {
  "type" : "string",
  "metadata" : {
    "description" : "Ignition content for the worker nodes"
  }
},
"numberOfNodes" : {
  "type" : "int",
  "defaultValue" : 3,
  "minValue" : 2,
  "maxValue" : 30,
  "metadata" : {
    "description" : "Number of OpenShift compute nodes to deploy"
  }
},
"sshKeyData" : {
  "type" : "securestring",
  "defaultValue" : "Unused",
  "metadata" : {
    "description" : "Unused"
  }
},
"nodeVMSize" : {
  "type" : "string",
  "defaultValue" : "Standard_D4s_v3",
  "metadata" : {
    "description" : "The size of the each Node Virtual Machine"
  }
},
"hyperVGen": {
  "type": "string",
  "metadata": {
    "description": "VM generation image to use"
  },
  "defaultValue": "V2",
  "allowedValues": [
    "V1",
    "V2"
  ]
},
"variables" : {
  "location" : "[resourceGroup().location]",
  "virtualNetworkName" : "[concat(if(not(empty(parameters('vnetBaseName'))),
parameters('vnetBaseName'), parameters('baseName')), '-vnet')]",
  "virtualNetworkID" : "[resourceId('Microsoft.Network/virtualNetworks',

```

```


variables('virtualNetworkName'))]",
  "nodeSubnetName" : "[concat(if(not(empty(parameters('vnetBaseName'))),
parameters('vnetBaseName'), parameters('baseName')), '-worker-subnet')]",
  "nodeSubnetRef" : "[concat(variables('virtualNetworkID'), '/subnets/',
variables('nodeSubnetName'))]",
  "infraLoadBalancerName" : "[parameters('baseName')]",
  "sshKeyPath" : "/home/capi/.ssh/authorized_keys",
  "identityName" : "[concat(parameters('baseName'), '-identity')]",
  "galleryName": "[concat('gallery_', replace(parameters('baseName'), '-', '_'))]",
  "imageName" : "[concat(parameters('baseName'), if(equals(parameters('hyperVGen'), 'V2'), '-
gen2', ''))]",
  "copy" : [
    {
      "name" : "vmNames",
      "count" : "[parameters('numberOfNodes')]",
      "input" : "[concat(parameters('baseName'), '-worker-', variables('location'), '-',
copyIndex('vmNames', 1))]"
    }
  ]
},
"resources" : [
  {
    "apiVersion" : "2019-05-01",
    "name" : "[concat('node', copyIndex())]",
    "type" : "Microsoft.Resources/deployments",
    "copy" : {
      "name" : "nodeCopy",
      "count" : "[length(variables('vmNames'))]"
    },
    "properties" : {
      "mode" : "Incremental",
      "template" : {
        "$schema" : "http://schema.management.azure.com/schemas/2015-01-
01/deploymentTemplate.json#",
        "contentVersion" : "1.0.0.0",
        "resources" : [
          {
            "apiVersion" : "2018-06-01",
            "type" : "Microsoft.Network/networkInterfaces",
            "name" : "[concat(variables('vmNames')[copyIndex()], '-nic')]",
            "location" : "[variables('location')]",
            "properties" : {
              "ipConfigurations" : [
                {
                  "name" : "pipConfig",
                  "properties" : {
                    "privateIPAllocationMethod" : "Dynamic",
                    "subnet" : {
                      "id" : "[variables('nodeSubnetRef')]"
                    }
                  }
                }
              ]
            }
          }
        ]
      }
    }
  },
  {

```

```

"apiVersion" : "2018-06-01",
"type" : "Microsoft.Compute/virtualMachines",
"name" : "[variables('vmNames')[copyIndex()]]",
"location" : "[variables('location')]",
"tags" : {
  "kubernetes.io-cluster-ffranzupi": "owned"
},
"identity" : {
  "type" : "userAssigned",
  "userAssignedIdentities" : {
    "[resourceID('Microsoft.ManagedIdentity/userAssignedIdentities/',
variables('identityName'))]" : {}
  }
},
"dependsOn" : [
  "[concat('Microsoft.Network/networkInterfaces/', concat(variables('vmNames')
[copyIndex()], '-nic'))]"
],
"properties" : {
  "hardwareProfile" : {
    "vmSize" : "[parameters('nodeVMSize')]"
  },
  "osProfile" : {
    "computerName" : "[variables('vmNames')[copyIndex()]]",
    "adminUsername" : "capi",
    "adminPassword" : "NotActuallyApplied!",
    "customData" : "[parameters('workerIgnition')]",
    "linuxConfiguration" : {
      "disablePasswordAuthentication" : false
    }
  },
  "storageProfile" : {
    "imageReference": {
      "id": "[resourceId('Microsoft.Compute/galleries/images', variables('galleryName'),
variables('imageName'))]"
    },
    "osDisk" : {
      "name": "[concat(variables('vmNames')[copyIndex()], '_OSDisk')]",
      "osType" : "Linux",
      "createOption" : "FromImage",
      "managedDisk": {
        "storageAccountType": "Premium_LRS"
      },
      "diskSizeGB": 128
    }
  },
  "networkProfile" : {
    "networkInterfaces" : [
      {
        "id" : "[resourceId('Microsoft.Network/networkInterfaces',
concat(variables('vmNames')[copyIndex()], '-nic'))]",
        "properties": {
          "primary": true
        }
      }
    ]
  }
}

```



7.11.20. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.



IMPORTANT

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**. To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:


```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.



NOTE

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your PATH.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

7.11.21. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

7.11.22. Approving the certificate signing requests for your machines

When you add machines to a cluster, two pending certificate signing requests (CSRs) are generated for each machine that you added. You must confirm that these CSRs are approved or, if necessary, approve them yourself. The client requests must be approved first, followed by the server requests.

Prerequisites

- You added machines to your cluster.

Procedure

1. Confirm that the cluster recognizes the machines:

```
$ oc get nodes
```

Example output

```
NAME      STATUS  ROLES  AGE  VERSION
master-0  Ready   master 63m  v1.29.4
master-1  Ready   master 63m  v1.29.4
master-2  Ready   master 64m  v1.29.4
```

The output lists all of the machines that you created.



NOTE

The preceding output might not include the compute nodes, also known as worker nodes, until some CSRs are approved.

- Review the pending CSRs and ensure that you see the client requests with the **Pending** or **Approved** status for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

```
NAME      AGE  REQUESTOR                                     CONDITION
csr-8b2br 15m  system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper  Pending
csr-8vnps 15m  system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper  Pending
...
```

In this example, two machines are joining the cluster. You might see more approved CSRs in the list.

- If the CSRs were not approved, after all of the pending CSRs for the machines you added are in **Pending** status, approve the CSRs for your cluster machines:



NOTE

Because the CSRs rotate automatically, approve your CSRs within an hour of adding the machines to the cluster. If you do not approve them within an hour, the certificates will rotate, and more than two certificates will be present for each node. You must approve all of these certificates. After the client CSR is approved, the Kubelet creates a secondary CSR for the serving certificate, which requires manual approval. Then, subsequent serving certificate renewal requests are automatically approved by the **machine-approver** if the Kubelet requests a new certificate with identical parameters.



NOTE

For clusters running on platforms that are not machine API enabled, such as bare metal and other user-provisioned infrastructure, you must implement a method of automatically approving the kubelet serving certificate requests (CSRs). If a request is not approved, then the **oc exec**, **oc rsh**, and **oc logs** commands cannot succeed, because a serving certificate is required when the API server connects to the kubelet. Any operation that contacts the Kubelet endpoint requires this certificate approval to be in place. The method must watch for new CSRs, confirm that the CSR was submitted by the **node-bootstrapper** service account in the **system:node** or **system:admin** groups, and confirm the identity of the node.

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

1 **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{\n"}\n{{end}}\n{{end}}' | xargs --no-run-if-empty oc adm certificate approve
```



NOTE

Some Operators might not become available until some CSRs are approved.

- Now that your client requests are approved, you must review the server requests for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

```
NAME      AGE   REQUESTOR                                     CONDITION
csr-bfd72 5m26s system:node:ip-10-0-50-126.us-east-2.compute.internal
Pending
csr-c57lv 5m26s system:node:ip-10-0-95-157.us-east-2.compute.internal
Pending
...
```

- If the remaining CSRs are not approved, and are in the **Pending** status, approve the CSRs for your cluster machines:

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

1 **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{\n"}\n{{end}}\n{{end}}' | xargs oc adm certificate approve
```

- After all client and server CSRs have been approved, the machines have the **Ready** status. Verify this by running the following command:

```
$ oc get nodes
```

Example output

```
NAME      STATUS   ROLES    AGE   VERSION
master-0 Ready    master  73m  v1.29.4
```

```

master-1 Ready   master 73m v1.29.4
master-2 Ready   master 74m v1.29.4
worker-0 Ready   worker 11m v1.29.4
worker-1 Ready   worker 11m v1.29.4

```



NOTE

It can take a few minutes after approval of the server CSRs for the machines to transition to the **Ready** status.

Additional information

- For more information on CSRs, see [Certificate Signing Requests](#).

7.11.23. Adding the Ingress DNS records

If you removed the DNS Zone configuration when creating Kubernetes manifests and generating Ignition configs, you must manually create DNS records that point at the Ingress load balancer. You can create either a wildcard ***.apps.{baseDomain}**, or specific records. You can use A, CNAME, and other records per your requirements.

Prerequisites

- You deployed an OpenShift Container Platform cluster on Microsoft Azure by using infrastructure that you provisioned.
- Install the OpenShift CLI (**oc**).
- Install or update the [Azure CLI](#).

Procedure

1. Confirm the Ingress router has created a load balancer and populated the **EXTERNAL-IP** field:

```
$ oc -n openshift-ingress get service router-default
```

Example output

```

NAME          TYPE          CLUSTER-IP    EXTERNAL-IP    PORT(S)          AGE
router-default LoadBalancer  172.30.20.10  35.130.120.110 80:32288/TCP,443:31215/TCP 20

```

2. Export the Ingress router IP as a variable:

```
$ export PUBLIC_IP_ROUTER=`oc -n openshift-ingress get service router-default --no-headers | awk '{print $4}'`
```

3. Add a ***.apps** record to the public DNS zone.

- a. If you are adding this cluster to a new public zone, run:

```
$ az network dns record-set a add-record -g ${BASE_DOMAIN_RESOURCE_GROUP} -z ${CLUSTER_NAME}.${BASE_DOMAIN} -n *.apps -a ${PUBLIC_IP_ROUTER} --ttl 300
```

- b. If you are adding this cluster to an already existing public zone, run:

```
$ az network dns record-set a add-record -g ${BASE_DOMAIN_RESOURCE_GROUP} -z ${BASE_DOMAIN} -n *.apps.${CLUSTER_NAME} -a ${PUBLIC_IP_ROUTER} --ttl 300
```

4. Add a ***.apps** record to the private DNS zone:

- a. Create a ***.apps** record by using the following command:

```
$ az network private-dns record-set a create -g ${RESOURCE_GROUP} -z ${CLUSTER_NAME}.${BASE_DOMAIN} -n *.apps --ttl 300
```

- b. Add the ***.apps** record to the private DNS zone by using the following command:

```
$ az network private-dns record-set a add-record -g ${RESOURCE_GROUP} -z ${CLUSTER_NAME}.${BASE_DOMAIN} -n *.apps -a ${PUBLIC_IP_ROUTER}
```

If you prefer to add explicit domains instead of using a wildcard, you can create entries for each of the cluster's current routes:

```
$ oc get --all-namespaces -o jsonpath='{range .items[*]}{range .status.ingress[*]}{.host}{"\n"}{end}{end}' routes
```

Example output

```
oauth-openshift.apps.cluster.basedomain.com
console-openshift-console.apps.cluster.basedomain.com
downloads-openshift-console.apps.cluster.basedomain.com
alertmanager-main-openshift-monitoring.apps.cluster.basedomain.com
prometheus-k8s-openshift-monitoring.apps.cluster.basedomain.com
```

7.11.24. Completing an Azure installation on user-provisioned infrastructure

After you start the OpenShift Container Platform installation on Microsoft Azure user-provisioned infrastructure, you can monitor the cluster events until the cluster is ready.

Prerequisites

- Deploy the bootstrap machine for an OpenShift Container Platform cluster on user-provisioned Azure infrastructure.
- Install the **oc** CLI and log in.

Procedure

- Complete the cluster installation:

```
$ ./openshift-install --dir <installation_directory> wait-for install-complete 1
```

Example output

```
INFO Waiting up to 30m0s for the cluster to initialize...
```

- 1 For `<installation_directory>`, specify the path to the directory that you stored the installation files in.



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

7.11.25. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service

7.12. INSTALLING A CLUSTER ON AZURE IN A RESTRICTED NETWORK

In OpenShift Container Platform version 4.16, you can install a cluster on Microsoft Azure in a restricted network by creating an internal mirror of the installation release content on an existing Azure Virtual Network (VNet).



IMPORTANT

You can install an OpenShift Container Platform cluster by using mirrored installation release content, but your cluster requires internet access to use the Azure APIs.

7.12.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).

- You [configured an Azure account](#) to host the cluster and determined the tested and validated region to deploy the cluster.
- You [mirrored the images for a disconnected installation](#) to your registry and obtained the `imageContentSources` data for your version of OpenShift Container Platform.



IMPORTANT

Because the installation media is on the mirror host, you can use that computer to complete all installation steps.

- You have an existing VNet in Azure. While installing a cluster in a restricted network that uses installer-provisioned infrastructure, you cannot use the installer-provisioned VNet. You must use a user-provisioned VNet that satisfies one of the following requirements:
 - The VNet contains the mirror registry
 - The VNet has firewall rules or a peering connection to access the mirror registry hosted elsewhere
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.
- If you use customer-managed encryption keys, you [prepared your Azure environment for encryption](#).

7.12.2. About installations in restricted networks

In OpenShift Container Platform 4.16, you can perform an installation that does not require an active connection to the internet to obtain software components. Restricted network installations can be completed using installer-provisioned infrastructure or user-provisioned infrastructure, depending on the cloud platform to which you are installing the cluster.

If you choose to perform a restricted network installation on a cloud platform, you still require access to its cloud APIs. Some cloud functions, like Amazon Web Service's Route 53 DNS and IAM services, require internet access. Depending on your network, you might require less internet access for an installation on bare metal hardware, Nutanix, or on VMware vSphere.

To complete a restricted network installation, you must create a registry that mirrors the contents of the OpenShift image registry and contains the installation media. You can create this registry on a mirror host, which can access both the internet and your closed network, or by using other methods that meet your restrictions.

7.12.2.1. Additional limits

Clusters in restricted networks have the following additional limitations and restrictions:

- The `ClusterVersion` status includes an **Unable to retrieve available updates** error.
- By default, you cannot use the contents of the Developer Catalog because you cannot access the required image stream tags.

7.12.2.2. User-defined outbound routing

In OpenShift Container Platform, you can choose your own outbound routing for a cluster to connect to the internet. This allows you to skip the creation of public IP addresses and the public load balancer.

You can configure user-defined routing by modifying parameters in the **install-config.yaml** file before installing your cluster. A pre-existing VNet is required to use outbound routing when installing a cluster; the installation program is not responsible for configuring this.

When configuring a cluster to use user-defined routing, the installation program does not create the following resources:

- Outbound rules for access to the internet.
- Public IPs for the public load balancer.
- Kubernetes Service object to add the cluster machines to the public load balancer for outbound requests.


You must ensure the following items are available before setting user-defined routing:

- Egress to the internet is possible to pull container images, unless using an OpenShift image registry mirror.
- The cluster can access Azure APIs.
- Various allowlist endpoints are configured. You can reference these endpoints in the *Configuring your firewall* section.

There are several pre-existing networking setups that are supported for internet access using user-defined routing.

Restricted cluster with Azure Firewall

You can use Azure Firewall to restrict the outbound routing for the Virtual Network (VNet) that is used to install the OpenShift Container Platform cluster. For more information, see [providing user-defined routing with Azure Firewall](#). You can create a OpenShift Container Platform cluster in a restricted network by using VNet with Azure Firewall and configuring the user-defined routing.



IMPORTANT

If you are using Azure Firewall for restricting internet access, you must set the **publish** field to **Internal** in the **install-config.yaml** file. This is because [Azure Firewall does not work properly with Azure public load balancers](#).

7.12.3. About reusing a VNet for your OpenShift Container Platform cluster

In OpenShift Container Platform 4.16, you can deploy a cluster into an existing Azure Virtual Network (VNet) in Microsoft Azure. If you do, you must also use existing subnets within the VNet and routing rules.

By deploying OpenShift Container Platform into an existing Azure VNet, you might be able to avoid service limit constraints in new accounts or more easily abide by the operational constraints that your company's guidelines set. This is a good option to use if you cannot obtain the infrastructure creation permissions that are required to create the VNet.

7.12.3.1. Requirements for using your VNet

When you deploy a cluster by using an existing VNet, you must perform additional network configuration before you install the cluster. In installer-provisioned infrastructure clusters, the installer usually creates the following components, but it does not create them when you install into an existing VNet:

- Subnets
- Route tables
- VNets
- Network Security Groups



NOTE

The installation program requires that you use the cloud-provided DNS server. Using a custom DNS server is not supported and causes the installation to fail.

If you use a custom VNet, you must correctly configure it and its subnets for the installation program and the cluster to use. The installation program cannot subdivide network ranges for the cluster to use, set route tables for the subnets, or set VNet options like DHCP, so you must do so before you install the cluster.

The cluster must be able to access the resource group that contains the existing VNet and subnets. While all of the resources that the cluster creates are placed in a separate resource group that it creates, some network resources are used from a separate group. Some cluster Operators must be able to access resources in both resource groups. For example, the Machine API controller attaches NICS for the virtual machines that it creates to subnets from the networking resource group.

Your VNet must meet the following characteristics:

- The VNet's CIDR block must contain the **Networking.MachineCIDR** range, which is the IP address pool for cluster machines.
- The VNet and its subnets must belong to the same resource group, and the subnets must be configured to use Azure-assigned DHCP IP addresses instead of static IP addresses.

You must provide two subnets within your VNet, one for the control plane machines and one for the compute machines. Because Azure distributes machines in different availability zones within the region that you specify, your cluster will have high availability by default.



NOTE

By default, if you specify availability zones in the **install-config.yaml** file, the installation program distributes the control plane machines and the compute machines across [these availability zones](#) within [a region](#). To ensure high availability for your cluster, select a region with at least three availability zones. If your region contains fewer than three availability zones, the installation program places more than one control plane machine in the available zones.

To ensure that the subnets that you provide are suitable, the installation program confirms the following data:

- All the specified subnets exist.
- There are two private subnets, one for the control plane machines and one for the compute machines.
- The subnet CIDRs belong to the machine CIDR that you specified. Machines are not provisioned in availability zones that you do not provide private subnets for. If required, the installation

program creates public load balancers that manage the control plane and worker nodes, and Azure allocates a public IP address to them.




NOTE

If you destroy a cluster that uses an existing VNet, the VNet is not deleted.

7.12.3.1.1. Network security group requirements

The network security groups for the subnets that host the compute and control plane machines require specific access to ensure that the cluster communication is correct. You must create rules to allow access to the required cluster communication ports.




IMPORTANT

The network security group rules must be in place before you install the cluster. If you attempt to install a cluster without the required access, the installation program cannot reach the Azure APIs, and installation fails.

Table 7.30. Required ports

Port	Description	Control plane	Compute
80	Allows HTTP traffic		x
443	Allows HTTPS traffic		x
6443	Allows communication to the control plane machines	x	
22623	Allows internal communication to the machine config server for provisioning machines	x	
*	Allows connections to Azure APIs. You must set a Destination Service Tag to AzureCloud . ^[1]	x	x
*	Denies connections to the internet. You must set a Destination Service Tag to Internet . ^[1]	x	x

1. If you are using Azure Firewall to restrict the internet access, then [you can configure Azure Firewall to allow the Azure APIs](#). A network security group rule is not needed.



IMPORTANT

Currently, there is no supported way to block or restrict the machine config server endpoint. The machine config server must be exposed to the network so that newly-provisioned machines, which have no existing configuration or state, are able to fetch their configuration. In this model, the root of trust is the certificate signing requests (CSR) endpoint, which is where the kubelet sends its certificate signing request for approval to join the cluster. Because of this, machine configs should not be used to distribute sensitive information, such as secrets and certificates.

To ensure that the machine config server endpoints, ports 22623 and 22624, are secured in bare metal scenarios, customers must configure proper network policies.

Because cluster components do not modify the user-provided network security groups, which the Kubernetes controllers update, a pseudo-network security group is created for the Kubernetes controller to modify without impacting the rest of the environment.

Additional resources

- [About the OpenShift SDN network plugin](#)
- [Configuring your firewall](#)

7.12.3.2. Division of permissions

Starting with OpenShift Container Platform 4.3, you do not need all of the permissions that are required for an installation program-provisioned infrastructure cluster to deploy a cluster. This change mimics the division of permissions that you might have at your company: some individuals can create different resources in your clouds than others. For example, you might be able to create application-specific items, like instances, storage, and load balancers, but not networking-related components such as VNets, subnet, or ingress rules.

The Azure credentials that you use when you create your cluster do not need the networking permissions that are required to make VNets and core networking components within the VNet, such as subnets, routing tables, internet gateways, NAT, and VPN. You still need permission to make the application resources that the machines within the cluster require, such as load balancers, security groups, storage accounts, and nodes.

7.12.3.3. Isolation between clusters

Because the cluster is unable to modify network security groups in an existing subnet, there is no way to isolate clusters from each other on the VNet.

7.12.4. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to obtain the images that are necessary to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.

- Obtain the packages that are required to perform cluster updates.

7.12.5. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.



NOTE

You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.



NOTE

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

-

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.



NOTE

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the `ssh-agent` process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```



NOTE

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the `ssh-agent`:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

7.12.6. Creating the installation configuration file

You can customize the OpenShift Container Platform cluster you install on Microsoft Azure.

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster. For a restricted network installation, these files are on your mirror host.

- You have the **imageContentSources** values that were generated during mirror registry creation.
- You have obtained the contents of the certificate for your mirror registry.
- You have retrieved a Red Hat Enterprise Linux CoreOS (RHCOS) image and uploaded it to an accessible location.
- You have an Azure subscription ID and tenant ID.
- If you are installing the cluster using a service principal, you have its application ID and password.
- If you are installing the cluster using a system-assigned managed identity, you have enabled it on the virtual machine that you will run the installation program from.
- If you are installing the cluster using a user-assigned managed identity, you have met these prerequisites:
 - You have its client ID.
 - You have assigned it to the virtual machine that you will run the installation program from.

Procedure

1. Optional: If you have run the installation program on this computer before, and want to use an alternative service principal or managed identity, go to the `~/.azure/` directory and delete the **osServicePrincipal.json** configuration file.
Deleting this file prevents the installation program from automatically reusing subscription and authentication values from a previous installation.
2. Create the **install-config.yaml** file.
 - a. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install create install-config --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the directory name to store the files that the installation program creates.

When specifying the directory:

- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
 - Use an empty directory. Some installation assets, such as bootstrap X.509 certificates, have short expiration intervals, therefore you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.
- b. At the prompts, provide the configuration details for your cloud:
 - i. Optional: Select an SSH key to use to access your cluster machines.

**NOTE**

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- ii. Select **azure** as the platform to target.
If the installation program cannot locate the **osServicePrincipal.json** configuration file from a previous installation, you are prompted for Azure subscription and authentication values.
- iii. Enter the following Azure parameter values for your subscription:
 - **azure subscription id** Enter the subscription ID to use for the cluster.
 - **azure tenant id** Enter the tenant ID.
- iv. Depending on the Azure identity you are using to deploy the cluster, do one of the following when prompted for the **azure service principal client id**
 - If you are using a service principal, enter its application ID.
 - If you are using a system-assigned managed identity, leave this value blank.
 - If you are using a user-assigned managed identity, specify its client ID.
- v. Depending on the Azure identity you are using to deploy the cluster, do one of the following when prompted for the **azure service principal client secret**
 - If you are using a service principal, enter its password.
 - If you are using a system-assigned managed identity, leave this value blank.
 - If you are using a user-assigned managed identity, leave this value blank.
- vi. Select the region to deploy the cluster to.
- vii. Select the base domain to deploy the cluster to. The base domain corresponds to the Azure DNS Zone that you created for your cluster.
- viii. Enter a descriptive name for your cluster.

**IMPORTANT**

All Azure resources that are available through public endpoints are subject to resource name restrictions, and you cannot create resources that use certain terms. For a list of terms that Azure restricts, see [Resolve reserved resource name errors](#) in the Azure documentation.

- ix. Paste the [pull secret from Red Hat OpenShift Cluster Manager](#) .
3. Edit the **install-config.yaml** file to give the additional information that is required for an installation in a restricted network.
 - a. Update the **pullSecret** value to contain the authentication information for your registry:


By setting this option, you create an internal Ingress Controller and a private load balancer.



IMPORTANT

Azure Firewall [does not work seamlessly](#) with Azure Public Load balancers. Thus, when using Azure Firewall for restricting internet access, the **publish** field in **install-config.yaml** should be set to **Internal**.

4. Make any other modifications to the **install-config.yaml** file that you require. For more information about the parameters, see "Installation configuration parameters".
5. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

If previously not detected, the installation program creates an **osServicePrincipal.json** configuration file and stores this file in the `~/.azure/` directory on your computer. This ensures that the installation program can load the profile when it is creating an OpenShift Container Platform cluster on the target platform.

Additional resources

- [Installation configuration parameters for Azure](#)

7.12.6.1. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 7.31. Minimum resource requirements

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Bootstrap	RHCOS	4	16 GB	100 GB	300
Control plane	RHCOS	4	16 GB	100 GB	300
Compute	RHCOS, RHEL 8.6 and later [3]	2	8 GB	100 GB	300

1. One vCPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio: $(\text{threads per core} \times \text{cores}) \times \text{sockets} = \text{vCPUs}$.
2. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes which require a 10 ms

p99 fsync duration. Note that on many cloud platforms, storage size and IOPS scale together, so you might need to over-allocate storage volume to obtain sufficient performance.

3. As with all user-provisioned installations, if you choose to use RHEL compute machines in your cluster, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Use of RHEL 7 compute machines is deprecated and has been removed in OpenShift Container Platform 4.10 and later.




NOTE

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).



IMPORTANT

You are required to use Azure virtual machines that have the **premiumIO** parameter set to **true**.

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

7.12.6.2. Tested instance types for Azure

The following Microsoft Azure instance types have been tested with OpenShift Container Platform.

Example 7.91. Machine types based on 64-bit x86 architecture

- **c4.***
- **c5.***
- **c5a.***
- **i3.***
- **m4.***
- **m5.***
- **m5a.***
- **m6a.***
- **m6i.***

- r4.*
- r5.*
- r5a.*
- r6i.*
- t3.*
- t3a.*

7.12.6.3. Tested instance types for Azure on 64-bit ARM infrastructures

The following Microsoft Azure ARM64 instance types have been tested with OpenShift Container Platform.


Example 7.92. Machine types based on 64-bit ARM architecture

- c6g.*
- m6g.*

7.12.6.4. Enabling trusted launch for Azure VMs

You can enable two trusted launch features when installing your cluster on Azure: [secure boot](#) and [virtualized Trusted Platform Modules](#).

See the Azure documentation about [virtual machine sizes](#) to learn what sizes of virtual machines support these features.



IMPORTANT

Trusted launch is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

Prerequisites

- You have created an **install-config.yaml** file.

Procedure

- Use a text editor to edit the **install-config.yaml** file prior to deploying your cluster and add the following stanza:

```
controlPlane: 1
```

```

platform:
  azure:
    settings:
      securityType: TrustedLaunch 2
    trustedLaunch:
      uefiSettings:
        secureBoot: Enabled 3
        virtualizedTrustedPlatformModule: Enabled 4

```

- 1 Specify **controlPlane.platform.azure** or **compute.platform.azure** to enable trusted launch on only control plane or compute nodes respectively. Specify **platform.azure.defaultMachinePlatform** to enable trusted launch on all nodes.
- 2 Enable trusted launch features.
- 3 Enable secure boot. For more information, see the Azure documentation about [secure boot](#).
- 4 Enable the virtualized Trusted Platform Module. For more information, see the Azure documentation about [virtualized Trusted Platform Modules](#).

7.12.6.5. Enabling confidential VMs

You can enable confidential VMs when installing your cluster. You can enable confidential VMs for compute nodes, control plane nodes, or all nodes.



IMPORTANT

Using confidential VMs is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

You can use confidential VMs with the following VM sizes:

- DCasv5-series
- DCadsv5-series
- ECasv5-series
- ECadsv5-series



IMPORTANT

Confidential VMs are currently not supported on 64-bit ARM architectures.

Prerequisites

- You have created an **install-config.yaml** file.

Procedure


- Use a text editor to edit the **install-config.yaml** file prior to deploying your cluster and add the following stanza:

```
controlPlane: ❶
  platform:
    azure:
      settings:
        securityType: ConfidentialVM ❷
        confidentialVM:
          uefiSettings:
            secureBoot: Enabled ❸
            virtualizedTrustedPlatformModule: Enabled ❹
      osDisk:
        securityProfile:
          securityEncryptionType: VMGuestStateOnly ❺
```

- ❶ Specify **controlPlane.platform.azure** or **compute.platform.azure** to deploy confidential VMs on only control plane or compute nodes respectively. Specify **platform.azure.defaultMachinePlatform** to deploy confidential VMs on all nodes.
- ❷ Enable confidential VMs.
- ❸ Enable secure boot. For more information, see the Azure documentation about [secure boot](#).
- ❹ Enable the virtualized Trusted Platform Module. For more information, see the Azure documentation about [virtualized Trusted Platform Modules](#).
- ❺ Specify **VMGuestStateOnly** to encrypt the VM guest state.

7.12.6.6. Sample customized install-config.yaml file for Azure

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.



IMPORTANT

This sample YAML file is provided for reference only. You must obtain your **install-config.yaml** file by using the installation program and modify it.

```
apiVersion: v1
baseDomain: example.com ❶
controlPlane: ❷
  hyperthreading: Enabled ❸ ❹
  name: master
  platform:
    azure:
      encryptionAtHost: true
      ultraSSDCapability: Enabled
      osDisk:
        diskSizeGB: 1024 ❺
```

```
diskType: Premium_LRS
diskEncryptionSet:
  resourceGroup: disk_encryption_set_resource_group
  name: disk_encryption_set_name
  subscriptionId: secondary_subscription_id
osImage:
  publisher: example_publisher_name
  offer: example_image_offer
  sku: example_offer_sku
  version: example_image_version
  type: Standard_D8s_v3
replicas: 3
compute: 6
- hyperthreading: Enabled 7
  name: worker
  platform:
    azure:
      ultraSSDCapability: Enabled
      type: Standard_D2s_v3
      encryptionAtHost: true
      osDisk:
        diskSizeGB: 512 8
        diskType: Standard_LRS
        diskEncryptionSet:
          resourceGroup: disk_encryption_set_resource_group
          name: disk_encryption_set_name
          subscriptionId: secondary_subscription_id
      osImage:
        publisher: example_publisher_name
        offer: example_image_offer
        sku: example_offer_sku
        version: example_image_version
      zones: 9
      - "1"
      - "2"
      - "3"
    replicas: 5
  metadata:
    name: test-cluster 10
  networking:
    clusterNetwork:
      - cidr: 10.128.0.0/14
      hostPrefix: 23
    machineNetwork:
      - cidr: 10.0.0.0/16
    networkType: OVNKubernetes 11
    serviceNetwork:
      - 172.30.0.0/16
  platform:
    azure:
      defaultMachinePlatform:
        osImage: 12
        publisher: example_publisher_name
        offer: example_image_offer
        sku: example_offer_sku
```

```

    version: example_image_version
    ultraSSDCapability: Enabled
    baseDomainResourceGroupName: resource_group 13
    region: centralus 14
    resourceGroupName: existing_resource_group 15
    networkResourceGroupName: vnet_resource_group 16
    virtualNetwork: vnet 17
    controlPlaneSubnet: control_plane_subnet 18
    computeSubnet: compute_subnet 19
    outboundType: UserDefinedRouting 20
    cloudName: AzurePublicCloud
    pullSecret: '{"auths": ...}' 21
    fips: false 22
    sshKey: ssh-ed25519 AAAA... 23
    additionalTrustBundle: | 24
      -----BEGIN CERTIFICATE-----
      <MY_TRUSTED_CA_CERT>
      -----END CERTIFICATE-----
    imageContentSources: 25
    - mirrors:
      - <local_registry>/<local_repository_name>/release
        source: quay.io/openshift-release-dev/ocp-release
    - mirrors:
      - <local_registry>/<local_repository_name>/release
        source: quay.io/openshift-release-dev/ocp-v4.0-art-dev
    publish: Internal 26

```

- 1 10 14 21 Required. The installation program prompts you for this value.
- 2 6 If you do not provide these parameters and values, the installation program provides the default value.
- 3 7 The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.
- 4 Whether to enable or disable simultaneous multithreading, or **hyperthreading**. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores. You can disable it by setting the parameter value to **Disabled**. If you disable simultaneous multithreading in some cluster machines, you must disable it in all cluster machines.



IMPORTANT

If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance. Use larger virtual machine types, such as **Standard_D8s_v3**, for your machines if you disable simultaneous multithreading.

- 5 8 You can specify the size of the disk to use in GB. Minimum recommendation for control plane nodes is 1024 GB.
- 9 Specify a list of zones to deploy your machines to. For high availability, specify at least two zones.

- 11 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 12 Optional: A custom Red Hat Enterprise Linux CoreOS (RHCOS) image that should be used to boot control plane and compute machines. The **publisher**, **offer**, **sku**, and **version** parameters under **platform.azure.defaultMachinePlatform.osImage** apply to both control plane and compute machines. If the parameters under **controlPlane.platform.azure.osImage** or **compute.platform.azure.osImage** are set, they override the **platform.azure.defaultMachinePlatform.osImage** parameters.
- 13 Specify the name of the resource group that contains the DNS zone for your base domain.
- 15 Specify the name of an already existing resource group to install your cluster to. If undefined, a new resource group is created for the cluster.
- 16 If you use an existing VNet, specify the name of the resource group that contains it.
- 17 If you use an existing VNet, specify its name.
- 18 If you use an existing VNet, specify the name of the subnet to host the control plane machines.
- 19 If you use an existing VNet, specify the name of the subnet to host the compute machines.
- 20 When using Azure Firewall to restrict Internet access, you must configure outbound routing to send traffic through the Azure Firewall. Configuring user-defined routing prevents exposing external endpoints in your cluster.
- 22 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.



IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 23 You can optionally provide the **sshKey** value that you use to access the machines in your cluster.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- 24 Provide the contents of the certificate file that you used for your mirror registry.

- 25 Provide the **imageContentSources** section from the output of the command to mirror the repository.
- 26 How to publish the user-facing endpoints of your cluster. When using Azure Firewall to restrict Internet access, set **publish** to **Internal** to deploy a private cluster. The user-facing endpoints then cannot be accessed from the internet. The default value is **External**.

7.12.6.7. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.

- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with `.` to match subdomains only. For
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.

**NOTE**

The installation program does not support the proxy **readinessEndpoints** field.

**NOTE**

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.

**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

7.12.7. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.

**IMPORTANT**

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**.
To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:


```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.



NOTE

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your PATH.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

7.12.8. Alternatives to storing administrator-level secrets in the kube-system project

By default, administrator secrets are stored in the **kube-system** project. If you configured the **credentialsMode** parameter in the **install-config.yaml** file to **Manual**, you must use one of the following alternatives:

- To manage long-term cloud credentials manually, follow the procedure in [Manually creating long-term credentials](#).
- To implement short-term credentials that are managed outside the cluster for individual components, follow the procedures in [Configuring an Azure cluster to use short-term credentials](#).

7.12.8.1. Manually creating long-term credentials

The Cloud Credential Operator (CCO) can be put into manual mode prior to installation in environments where the cloud identity and access management (IAM) APIs are not reachable, or the administrator prefers not to store an administrator-level credential secret in the cluster **kube-system** namespace.

Procedure

1. If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```
-
```

```

apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...

```

- If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

- Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

- Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```

$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included 1 \
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml 2 \
  --to=<path_to_directory_for_credentials_requests> 3

```

- The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- Specify the location of the **install-config.yaml** file.
- Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

This command creates a YAML file for each **CredentialsRequest** object.

Sample **CredentialsRequest** object

```

apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: AzureProviderSpec
    roleBindings:
      - role: Contributor
    ...

```

5. Create YAML files for secrets in the **openshift-install** manifests directory that you generated previously. The secrets must be stored using the namespace and secret name defined in the **spec.secretRef** for each **CredentialsRequest** object.

Sample CredentialsRequest object with secrets

```

apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: AzureProviderSpec
    roleBindings:
      - role: Contributor
      ...
  secretRef:
    name: <component_secret>
    namespace: <component_namespace>
    ...

```

Sample Secret object

```

apiVersion: v1
kind: Secret
metadata:
  name: <component_secret>
  namespace: <component_namespace>
data:
  azure_subscription_id: <base64_encoded_azure_subscription_id>
  azure_client_id: <base64_encoded_azure_client_id>
  azure_client_secret: <base64_encoded_azure_client_secret>
  azure_tenant_id: <base64_encoded_azure_tenant_id>
  azure_resource_prefix: <base64_encoded_azure_resource_prefix>
  azure_resourcegroup: <base64_encoded_azure_resourcegroup>
  azure_region: <base64_encoded_azure_region>

```



IMPORTANT

Before upgrading a cluster that uses manually maintained credentials, you must ensure that the CCO is in an upgradeable state.

7.12.8.2. Configuring an Azure cluster to use short-term credentials

To install a cluster that uses Microsoft Entra Workload ID, you must configure the Cloud Credential Operator utility and create the required Azure resources for your cluster.

7.12.8.2.1. Configuring the Cloud Credential Operator utility

To create and manage cloud credentials from outside of the cluster when the Cloud Credential Operator (CCO) is operating in manual mode, extract and prepare the CCO utility (**ccoctl**) binary.



NOTE

The **ccoctl** utility is a Linux binary that must run in a Linux environment.

Prerequisites

- You have access to an OpenShift Container Platform account with cluster administrator access.
- You have installed the OpenShift CLI (**oc**).
- You have created a global Microsoft Azure account for the **ccoctl** utility to use with the following permissions:

Example 7.93. Required Azure permissions

- Microsoft.Resources/subscriptions/resourceGroups/read
- Microsoft.Resources/subscriptions/resourceGroups/write
- Microsoft.Resources/subscriptions/resourceGroups/delete
- Microsoft.Authorization/roleAssignments/read
- Microsoft.Authorization/roleAssignments/delete
- Microsoft.Authorization/roleAssignments/write
- Microsoft.Authorization/roleDefinitions/read
- Microsoft.Authorization/roleDefinitions/write
- Microsoft.Authorization/roleDefinitions/delete
- Microsoft.Storage/storageAccounts/listkeys/action
- Microsoft.Storage/storageAccounts/delete
- Microsoft.Storage/storageAccounts/read
- Microsoft.Storage/storageAccounts/write
- Microsoft.Storage/storageAccounts/blobServices/containers/write
- Microsoft.Storage/storageAccounts/blobServices/containers/delete
- Microsoft.Storage/storageAccounts/blobServices/containers/read
- Microsoft.ManagedIdentity/userAssignedIdentities/delete
- Microsoft.ManagedIdentity/userAssignedIdentities/read
- Microsoft.ManagedIdentity/userAssignedIdentities/write
- Microsoft.ManagedIdentity/userAssignedIdentities/federatedIdentityCredentials/read
- Microsoft.ManagedIdentity/userAssignedIdentities/federatedIdentityCredentials/write
- Microsoft.ManagedIdentity/userAssignedIdentities/federatedIdentityCredentials/delete

- Microsoft.Storage/register/action
- Microsoft.ManagedIdentity/register/action

Procedure

1. Set a variable for the OpenShift Container Platform release image by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

2. Obtain the CCO container image from the OpenShift Container Platform release image by running the following command:

```
$ CCO_IMAGE=$(oc adm release info --image-for='cloud-credential-operator'
$RELEASE_IMAGE -a ~/.pull-secret)
```



NOTE

Ensure that the architecture of the **\$RELEASE_IMAGE** matches the architecture of the environment in which you will use the **ccoctl** tool.

3. Extract the **ccoctl** binary from the CCO container image within the OpenShift Container Platform release image by running the following command:

```
$ oc image extract $CCO_IMAGE \
--file="/usr/bin/ccoctl.<rhel_version>" \
-a ~/.pull-secret
```

- 1 For **<rhel_version>**, specify the value that corresponds to the version of Red Hat Enterprise Linux (RHEL) that the host uses. If no value is specified, **ccoctl.rhel8** is used by default. The following values are valid:

- **rhel8**: Specify this value for hosts that use RHEL 8.
- **rhel9**: Specify this value for hosts that use RHEL 9.

4. Change the permissions to make **ccoctl** executable by running the following command:

```
$ chmod 775 ccoctl.<rhel_version>
```

Verification

- To verify that **ccoctl** is ready to use, display the help file by running the following command:

```
$ ccoctl --help
```

Output of **ccoctl --help**

```
OpenShift credentials provisioning tool
```


Usage:
 ccoctl [command]

Available Commands:

aws Manage credentials objects for AWS cloud
 azure Manage credentials objects for Azure
 gcp Manage credentials objects for Google cloud
 help Help about any command
 ibmcloud Manage credentials objects for IBM Cloud
 nutanix Manage credentials objects for Nutanix

Flags:

-h, --help help for ccoctl

Use "ccoctl [command] --help" for more information about a command.

7.12.8.2.2. Creating Azure resources with the Cloud Credential Operator utility

You can use the **ccoctl azure create-all** command to automate the creation of Azure resources.



NOTE

By default, **ccoctl** creates objects in the directory in which the commands are run. To create the objects in a different directory, use the **--output-dir** flag. This procedure uses **<path_to_ccoctl_output_dir>** to refer to this directory.

Prerequisites

You must have:

- Extracted and prepared the **ccoctl** binary.
- Access to your Microsoft Azure account by using the Azure CLI.

Procedure

1. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

2. Extract the list of **CredentialsRequest** objects from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included 1 \
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml 2 \
  --to=<path_to_directory_for_credentials_requests> 3
```

- 1** The **--included** parameter includes only the manifests that your specific cluster configuration requires.

- 2 Specify the location of the **install-config.yaml** file.
- 3 Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.



NOTE

This command might take a few moments to run.

3. To enable the **ccoctl** utility to detect your Azure credentials automatically, log in to the Azure CLI by running the following command:

```
$ az login
```

4. Use the **ccoctl** tool to process all **CredentialsRequest** objects by running the following command:

```
$ ccoctl azure create-all \
  --name=<azure_infra_name> \ 1
  --output-dir=<ccoctl_output_dir> \ 2
  --region=<azure_region> \ 3
  --subscription-id=<azure_subscription_id> \ 4
  --credentials-requests-dir=<path_to_credentials_requests_directory> \ 5
  --dnszone-resource-group-name=<azure_dns_zone_resource_group_name> \ 6
  --tenant-id=<azure_tenant_id> \ 7
```

- 1 Specify the user-defined name for all created Azure resources used for tracking.
- 2 Optional: Specify the directory in which you want the **ccoctl** utility to create objects. By default, the utility creates objects in the directory in which the commands are run.
- 3 Specify the Azure region in which cloud resources will be created.
- 4 Specify the Azure subscription ID to use.
- 5 Specify the directory containing the files for the component **CredentialsRequest** objects.
- 6 Specify the name of the resource group containing the cluster's base domain Azure DNS zone.
- 7 Specify the Azure tenant ID to use.



NOTE

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

To see additional optional parameters and explanations of how to use them, run the **azure create-all --help** command.

Verification

- To verify that the OpenShift Container Platform secrets are created, list the files in the `<path_to_ccoctl_output_dir>/manifests` directory:

```
$ ls <path_to_ccoctl_output_dir>/manifests
```

Example output

```
azure-ad-pod-identity-webhook-config.yaml
cluster-authentication-02-config.yaml
openshift-cloud-controller-manager-azure-cloud-credentials-credentials.yaml
openshift-cloud-network-config-controller-cloud-credentials-credentials.yaml
openshift-cluster-api-capz-manager-bootstrap-credentials-credentials.yaml
openshift-cluster-csi-drivers-azure-disk-credentials-credentials.yaml
openshift-cluster-csi-drivers-azure-file-credentials-credentials.yaml
openshift-image-registry-installer-cloud-credentials-credentials.yaml
openshift-ingress-operator-cloud-credentials-credentials.yaml
openshift-machine-api-azure-cloud-credentials-credentials.yaml
```

You can verify that the Microsoft Entra ID service accounts are created by querying Azure. For more information, refer to Azure documentation on listing Entra ID service accounts.

7.12.8.2.3. Incorporating the Cloud Credential Operator utility manifests

To implement short-term security credentials managed outside the cluster for individual components, you must move the manifest files that the Cloud Credential Operator utility (**ccoctl**) created to the correct directories for the installation program.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have configured the Cloud Credential Operator utility (**ccoctl**).
- You have created the cloud provider resources that are required for your cluster with the **ccoctl** utility.

Procedure

1. If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```
apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...
```

2. If you used the **ccoctl** utility to create a new Azure resource group instead of using an existing resource group, modify the **resourceGroupName** parameter in the **install-config.yaml** as shown:

Sample configuration file snippet

```

apiVersion: v1
baseDomain: example.com
# ...
platform:
  azure:
    resourceGroupName: <azure_infra_name> 1
# ...

```

- 1** This value must match the user-defined name for Azure resources that was specified with the **--name** argument of the **ccoctl azure create-all** command.

3. If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

4. Copy the manifests that the **ccoctl** utility generated to the **manifests** directory that the installation program created by running the following command:

```
$ cp /<path_to_ccoctl_output_dir>/manifests/* ./manifests/
```

5. Copy the private key that the **ccoctl** utility generated in the **tls** directory to the installation directory by running the following command:

```
$ cp -a /<path_to_ccoctl_output_dir>/tls .
```

7.12.9. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have an Azure subscription ID and tenant ID.

Procedure

- Change to the directory that contains the installation program and initialize the cluster deployment:


```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

- 1 For **<installation_directory>**, specify the location of your customized **./install-config.yaml** file.
- 2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to **<installation_directory>/openshift_install.log**.




IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

7.12.10. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the

correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

7.12.11. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service

7.12.12. Next steps

- [Customize your cluster](#).
- If necessary, you can [opt out of remote health reporting](#) .

7.13. INSTALLING A THREE-NODE CLUSTER ON AZURE

In OpenShift Container Platform version 4.16, you can install a three-node cluster on Microsoft Azure. A three-node cluster consists of three control plane machines, which also act as compute machines. This type of cluster provides a smaller, more resource efficient cluster, for cluster administrators and developers to use for testing, development, and production.

You can install a three-node cluster using either installer-provisioned or user-provisioned infrastructure.



NOTE

Deploying a three-node cluster using an Azure Marketplace image is not supported.

7.13.1. Configuring a three-node cluster

You configure a three-node cluster by setting the number of worker nodes to **0** in the **install-config.yaml** file before deploying the cluster. Setting the number of worker nodes to **0** ensures that the control plane machines are schedulable. This allows application workloads to be scheduled to run from the control plane nodes.



NOTE

Because application workloads run from control plane nodes, additional subscriptions are required, as the control plane nodes are considered to be compute nodes.

Prerequisites

- You have an existing **install-config.yaml** file.

Procedure

1. Set the number of compute replicas to **0** in your **install-config.yaml** file, as shown in the following **compute** stanza:

Example **install-config.yaml** file for a three-node cluster

```
apiVersion: v1
baseDomain: example.com
compute:
- name: worker
  platform: {}
  replicas: 0
# ...
```

2. If you are deploying a cluster with user-provisioned infrastructure:
 - After you create the Kubernetes manifest files, make sure that the **spec.mastersSchedulable** parameter is set to **true** in **cluster-scheduler-02-config.yml** file. You can locate this file in **<installation_directory>/manifests**. For more information, see "Creating the Kubernetes manifest and Ignition config files" in "Installing a cluster on Azure using ARM templates".
 - Do not create additional worker nodes.

Example **cluster-scheduler-02-config.yml** file for a three-node cluster

```
apiVersion: config.openshift.io/v1
kind: Scheduler
metadata:
  creationTimestamp: null
  name: cluster
```

```
spec:
  mastersSchedulable: true
  policy:
    name: ""
  status: {}
```

7.13.2. Next steps

- [Installing a cluster on Azure with customizations](#)
- [Installing a cluster on Azure using ARM templates](#)

7.14. UNINSTALLING A CLUSTER ON AZURE

You can remove a cluster that you deployed to Microsoft Azure.

7.14.1. Removing a cluster that uses installer-provisioned infrastructure

You can remove a cluster that uses installer-provisioned infrastructure from your cloud.



NOTE

After uninstallation, check your cloud provider for any resources not removed properly, especially with User Provisioned Infrastructure (UPI) clusters. There might be resources that the installer did not create or that the installer is unable to access.

Prerequisites

- You have a copy of the installation program that you used to deploy the cluster.
- You have the files that the installation program generated when you created your cluster.

Procedure

1. From the directory that contains the installation program on the computer that you used to install the cluster, run the following command:

```
$ ./openshift-install destroy cluster \
--dir <installation_directory> --log-level info 1 2
```

1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2 To view different details, specify **warn**, **debug**, or **error** instead of **info**.



NOTE

You must specify the directory that contains the cluster definition files for your cluster. The installation program requires the **metadata.json** file in this directory to delete the cluster.

- Optional: Delete the **<installation_directory>** directory and the OpenShift Container Platform installation program.

7.14.2. Deleting Microsoft Azure resources with the Cloud Credential Operator utility

After uninstalling an OpenShift Container Platform cluster that uses short-term credentials managed outside the cluster, you can use the CCO utility (**ccoctl**) to remove the Microsoft Azure (Azure) resources that **ccoctl** created during installation.

Prerequisites

- Extract and prepare the **ccoctl** binary.
- Uninstall an OpenShift Container Platform cluster on Azure that uses short-term credentials.

Procedure

- Delete the Azure resources that **ccoctl** created by running the following command:

```
$ ccoctl azure delete \
  --name=<name> \ 1
  --region=<azure_region> \ 2
  --subscription-id=<azure_subscription_id> \ 3
  --delete-oidc-resource-group
```

- 1** **<name>** matches the name that was originally used to create and tag the cloud resources.
- 2** **<azure_region>** is the Azure region in which to delete cloud resources.
- 3** **<azure_subscription_id>** is the Azure subscription ID for which to delete cloud resources.

Verification

- To verify that the resources are deleted, query Azure. For more information, refer to Azure documentation.

7.15. INSTALLATION CONFIGURATION PARAMETERS FOR AZURE

Before you deploy an OpenShift Container Platform cluster on Microsoft Azure, you provide parameters to customize your cluster and the platform that hosts it. When you create the **install-config.yaml** file, you provide values for the required parameters through the command line. You can then modify the **install-config.yaml** file to customize your cluster further.

7.15.1. Available installation configuration parameters for Azure

The following tables specify the required, optional, and Azure-specific installation configuration parameters that you can set as part of the installation process.



NOTE

After installation, you cannot modify these parameters in the **install-config.yaml** file.

7.15.1.1. Required configuration parameters

Required installation configuration parameters are described in the following table:

Table 7.32. Required parameters

Parameter	Description	Values
<code>apiVersion:</code>	The API version for the install-config.yaml content. The current version is v1 . The installation program may also support older API versions.	String
<code>baseDomain:</code>	The base domain of your cloud provider. The base domain is used to create routes to your OpenShift Container Platform cluster components. The full DNS name for your cluster is a combination of the baseDomain and metadata.name parameter values that uses the <metadata.name> . <baseDomain> format.	A fully-qualified domain or subdomain name, such as example.com .
<code>metadata:</code>	Kubernetes resource ObjectMeta , from which only the name parameter is consumed.	Object
<code>metadata: name:</code>	The name of the cluster. DNS records for the cluster are all subdomains of {{.metadata.name}} . {{.baseDomain}} .	String of lowercase letters, hyphens (-), and periods (.), such as dev .
<code>platform:</code>	The configuration for the specific platform upon which to perform the installation: aws , baremetal , azure , gcp , ibmcloud , nutanix , openstack , powervs , vsphere , or {} . For additional information about platform . <platform> parameters, consult the table for your specific platform that follows.	Object

Parameter	Description	Values
<code>pullSecret:</code>	Get a pull secret from Red Hat OpenShift Cluster Manager to authenticate downloading container images for OpenShift Container Platform components from services such as Quay.io.	<pre>{ "auths":{ "cloud.openshift.com":{ "auth":"b3Blb=", "email":"you@example.com" }, "quay.io":{ "auth":"b3Blb=", "email":"you@example.com" } } }</pre>

7.15.1.2. Network configuration parameters

You can customize your installation configuration based on the requirements of your existing network infrastructure. For example, you can expand the IP address block for the cluster network or provide different IP address blocks than the defaults.

Only IPv4 addresses are supported.




NOTE

Globalnet is not supported with Red Hat OpenShift Data Foundation disaster recovery solutions. For regional disaster recovery scenarios, ensure that you use a nonoverlapping range of private IP addresses for the cluster and service networks in each cluster.

Table 7.33. Network parameters

Parameter	Description	Values
<code>networking:</code>	The configuration for the cluster network.	Object <div style="display: flex; align-items: center; margin-top: 10px;"> <div> <p>NOTE</p> <p>You cannot modify parameters specified by the networking object after installation.</p> </div> </div>
<code>networking: networkType:</code>	The Red Hat OpenShift Networking network plugin to install.	OVNKubernetes. OVNKubernetes is a CNI plugin for Linux networks and hybrid networks that contain both Linux and Windows servers. The default value is OVNKubernetes .


Parameter	Description	Values
<code>networking: clusterNetwork:</code>	<p>The IP address blocks for pods.</p> <p>The default value is 10.128.0.0/14 with a host prefix of /23.</p> <p>If you specify multiple IP address blocks, the blocks must not overlap.</p>	<p>An array of objects. For example:</p> <pre>networking: clusterNetwork: - cidr: 10.128.0.0/14 hostPrefix: 23</pre>
<code>networking: clusterNetwork: cidr:</code>	<p>Required if you use networking.clusterNetwork. An IP address block.</p> <p>An IPv4 network.</p>	<p>An IP address block in Classless Inter-Domain Routing (CIDR) notation. The prefix length for an IPv4 block is between 0 and 32.</p>
<code>networking: clusterNetwork: hostPrefix:</code>	<p>The subnet prefix length to assign to each individual node. For example, if hostPrefix is set to 23 then each node is assigned a /23 subnet out of the given cidr. A hostPrefix value of 23 provides $2^{(32 - 23)} - 2$ pod IP addresses.</p>	<p>A subnet prefix.</p> <p>The default value is 23.</p>
<code>networking: serviceNetwork:</code>	<p>The IP address block for services. The default value is 172.30.0.0/16.</p> <p>The OVN-Kubernetes network plugins supports only a single IP address block for the service network.</p>	<p>An array with an IP address block in CIDR format. For example:</p> <pre>networking: serviceNetwork: - 172.30.0.0/16</pre>
<code>networking: machineNetwork:</code>	<p>The IP address blocks for machines.</p> <p>If you specify multiple IP address blocks, the blocks must not overlap.</p>	<p>An array of objects. For example:</p> <pre>networking: machineNetwork: - cidr: 10.0.0.0/16</pre>
<code>networking: machineNetwork: cidr:</code>	<p>Required if you use networking.machineNetwork. An IP address block. The default value is 10.0.0.0/16 for all platforms other than libvirt and IBM Power® Virtual Server. For libvirt, the default value is 192.168.126.0/24. For IBM Power® Virtual Server, the default value is 192.168.0.0/24.</p>	<p>An IP network block in CIDR notation.</p> <p>For example, 10.0.0.0/16.</p> <div style="display: flex; align-items: flex-start;">  <div> <p>NOTE</p> <p>Set the networking.machineNetwork to match the CIDR that the preferred NIC resides in.</p> </div> </div>


7.15.1.3. Optional configuration parameters

Optional installation configuration parameters are described in the following table:

Table 7.34. Optional parameters


Parameter	Description	Values
<code>additionalTrustBundle:</code>	A PEM-encoded X.509 certificate bundle that is added to the nodes' trusted certificate store. This trust bundle may also be used when a proxy has been configured.	String
<code>capabilities:</code>	Controls the installation of optional core cluster components. You can reduce the footprint of your OpenShift Container Platform cluster by disabling optional components. For more information, see the "Cluster capabilities" page in <i>Installing</i> .	String array
<code>capabilities: baselineCapabilitySet:</code>	Selects an initial set of optional capabilities to enable. Valid values are None , v4.11 , v4.12 and vCurrent . The default value is vCurrent .	String
<code>capabilities: additionalEnabledCapabilities:</code>	Extends the set of optional capabilities beyond what you specify in baselineCapabilitySet . You may specify multiple capabilities in this parameter.	String array
<code>cpuPartitioningMode:</code>	Enables workload partitioning, which isolates OpenShift Container Platform services, cluster management workloads, and infrastructure pods to run on a reserved set of CPUs. Workload partitioning can only be enabled during installation and cannot be disabled after installation. While this field enables workload partitioning, it does not configure workloads to use specific CPUs. For more information, see the <i>Workload partitioning</i> page in the <i>Scalability and Performance</i> section.	None or AllNodes . None is the default value.
<code>compute:</code>	The configuration for the machines that comprise the compute nodes.	Array of MachinePool objects.

Parameter	Description	Values
<code>compute: architecture:</code>	Determines the instruction set architecture of the machines in the pool. Currently, clusters with varied architectures are not supported. All pools must specify the same architecture. Valid values are amd64 and arm64 . Not all installation options support the 64-bit ARM architecture. To verify if your installation option is supported on your platform, see <i>Supported installation methods for different platforms</i> in <i>Selecting a cluster installation method and preparing it for users</i> .	String
<code>compute: hyperthreading:</code>	Whether to enable or disable simultaneous multithreading, or hyperthreading , on compute machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.  IMPORTANT If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.	Enabled or Disabled
<code>compute: name:</code>	Required if you use compute . The name of the machine pool.	worker
<code>compute: platform:</code>	Required if you use compute . Use this parameter to specify the cloud provider to host the worker machines. This parameter value must match the controlPlane.platform parameter value.	aws, azure, gcp, ibmcloud, nutanix, openstack, powervs, vsphere, or {}
<code>compute: replicas:</code>	The number of compute machines, which are also known as worker machines, to provision.	A positive integer greater than or equal to 2 . The default value is 3 .

Parameter	Description	Values
featureSet:	Enables the cluster for a feature set. A feature set is a collection of OpenShift Container Platform features that are not enabled by default. For more information about enabling a feature set during installation, see "Enabling features using feature gates".	String. The name of the feature set to enable, such as TechPreviewNoUpgrade .
controlPlane:	The configuration for the machines that comprise the control plane.	Array of MachinePool objects.
controlPlane: architecture:	Determines the instruction set architecture of the machines in the pool. Currently, clusters with varied architectures are not supported. All pools must specify the same architecture. Valid values are amd64 and arm64 . Not all installation options support the 64-bit ARM architecture. To verify if your installation option is supported on your platform, see <i>Supported installation methods for different platforms</i> in <i>Selecting a cluster installation method and preparing it for users</i> .	String
controlPlane: hyperthreading:	Whether to enable or disable simultaneous multithreading, or hyperthreading , on control plane machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.  IMPORTANT If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.	Enabled or Disabled

Parameter	Description	Values
<code>controlPlane: name:</code>	Required if you use controlPlane . The name of the machine pool.	master
<code>controlPlane: platform:</code>	Required if you use controlPlane . Use this parameter to specify the cloud provider that hosts the control plane machines. This parameter value must match the compute.platform parameter value.	aws, azure, gcp, ibmcloud, nutanix, openstack, powervs, vsphere , or {}
<code>controlPlane: replicas:</code>	The number of control plane machines to provision.	Supported values are 3 , or 1 when deploying single-node OpenShift.
<code>credentialsMode:</code>	The Cloud Credential Operator (CCO) mode. If no mode is specified, the CCO dynamically tries to determine the capabilities of the provided credentials, with a preference for mint mode on the platforms where multiple modes are supported.	Mint, Passthrough, Manual or an empty string (<code>""</code>). ^[1]

Parameter	Description	Values
<p>fips:</p>	<p>Enable or disable FIPS mode. The default is false (disabled). If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.</p> <div data-bbox="486 584 592 1637" style="background-color: black; width: 66px; height: 470px; margin-bottom: 10px;"></div> <p>IMPORTANT</p> <p>To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see Installing the system in FIPS mode.</p> <p>When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.</p> <div data-bbox="486 1682 592 1883" style="background-color: black; width: 66px; height: 90px; margin-bottom: 10px;"></div> <p>NOTE</p> <p>If you are using Azure File storage, you cannot enable FIPS mode.</p>	<p>false or true</p>
<p>imageContentSources:</p>	<p>Sources and repositories for the release-image content.</p>	<p>Array of objects. Includes a source and, optionally, mirrors, as described in the following rows of this table.</p>

Parameter	Description	Values
<code>imageContentSources:</code> <code>source:</code>	Required if you use imageContentSources . Specify the repository that users refer to, for example, in image pull specifications.	String
<code>imageContentSources:</code> <code>mirrors:</code>	Specify one or more repositories that may also contain the same images.	Array of strings
<code>publish:</code>	How to publish or expose the user-facing endpoints of your cluster, such as the Kubernetes API, OpenShift routes.	Internal , External , or Mixed . To deploy a private cluster, which cannot be accessed from the internet, set publish to Internal . The default value is External . To deploy a cluster where the API and the ingress server have different publishing strategies, set publish to Mixed and use the operatorPublishingStrategy parameter.
<code>sshKey:</code>	The SSH key to authenticate access to your cluster machines.  NOTE For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your ssh-agent process uses.	For example, sshKey: ssh-ed25519 AAAA...

1. Not all CCO modes are supported for all cloud providers. For more information about CCO modes, see the "Managing cloud provider credentials" entry in the *Authentication and authorization* content.



IMPORTANT

Setting this parameter to **Manual** enables alternatives to storing administrator-level secrets in the **kube-system** project, which require additional configuration steps. For more information, see "Alternatives to storing administrator-level secrets in the kube-system project".

7.15.1.4. Additional Azure configuration parameters

Additional Azure configuration parameters are described in the following table.



NOTE

By default, if you specify availability zones in the **install-config.yaml** file, the installation program distributes the control plane machines and the compute machines across [these availability zones](#) within [a region](#). To ensure high availability for your cluster, select a region with at least three availability zones. If your region contains fewer than three availability zones, the installation program places more than one control plane machine in the available zones.

Table 7.35. Additional Azure parameters

Parameter	Description	Values
<pre>compute: platform: azure: encryptionAtHost:</pre>	Enables host-level encryption for compute machines. You can enable this encryption alongside user-managed server-side encryption. This feature encrypts temporary, ephemeral, cached and un-managed disks on the VM host. This is not a prerequisite for user-managed server-side encryption.	true or false . The default is false .
<pre>compute: platform: azure: osDisk: diskSizeGB:</pre>	The Azure disk size for the VM.	Integer that represents the size of the disk in GB. The default is 128 .
<pre>compute: platform: azure: osDisk: diskType:</pre>	Defines the type of disk.	standard_LRS , premium_LRS , or standardSSD_LRS . The default is premium_LRS .
<pre>compute: platform: azure: ultraSSDCapability:</pre>	Enables the use of Azure ultra disks for persistent storage on compute nodes. This requires that your Azure region and zone have ultra disks available.	Enabled , Disabled . The default is Disabled .

Parameter	Description	Values
<pre>compute: platform: azure: osDisk: diskEncryptionSet: resourceGroup:</pre>	<p>The name of the Azure resource group that contains the disk encryption set from the installation prerequisites. This resource group should be different from the resource group where you install the cluster to avoid deleting your Azure encryption key when the cluster is destroyed. This value is only necessary if you intend to install the cluster with user-managed disk encryption.</p>	<p>String, for example production_encryption_resource_group.</p>
<pre>compute: platform: azure: osDisk: diskEncryptionSet: name:</pre>	<p>The name of the disk encryption set that contains the encryption key from the installation prerequisites.</p>	<p>String, for example production_disk_encryption_set.</p>
<pre>compute: platform: azure: osDisk: diskEncryptionSet: subscriptionId:</pre>	<p>Defines the Azure subscription of the disk encryption set where the disk encryption set resides. This secondary disk encryption set is used to encrypt compute machines.</p>	<p>String, in the format 00000000-0000-0000-0000-000000000000.</p>
<pre>compute: platform: azure: osImage: publisher:</pre>	<p>Optional. By default, the installation program downloads and installs the Red Hat Enterprise Linux CoreOS (RHCOS) image that is used to boot compute machines. You can override the default behavior by using a custom RHCOS image that is available from the Azure Marketplace. The installation program uses this image for compute machines only.</p>	<p>String. The name of the image publisher.</p>
<pre>compute: platform: azure: osImage: offer:</pre>	<p>The name of Azure Marketplace offer that is associated with the custom RHCOS image. If you use compute.platform.azure.osImage.publisher, this field is required.</p>	<p>String. The name of the image offer.</p>

Parameter	Description	Values
<pre>compute: platform: azure: osImage: sku:</pre>	<p>An instance of the Azure Marketplace offer. If you use compute.platform.azure.osImage.publisher, this field is required.</p>	String. The SKU of the image offer.
<pre>compute: platform: azure: osImage: version:</pre>	<p>The version number of the image SKU. If you use compute.platform.azure.osImage.publisher, this field is required.</p>	String. The version of the image to use.
<pre>compute: platform: azure: vmNetworkingType :</pre>	<p>Enables accelerated networking. Accelerated networking enables single root I/O virtualization (SR-IOV) to a VM, improving its networking performance. If instance type of compute machines support Accelerated networking, by default, the installer enables Accelerated networking, otherwise the default networking type is Basic.</p>	Accelerated or Basic .
<pre>compute: platform: azure: type:</pre>	<p>Defines the Azure instance type for compute machines.</p>	String
<pre>compute: platform: azure: zones:</pre>	<p>The availability zones where the installation program creates compute machines.</p>	String list
<pre>compute: platform: azure: settings: securityType:</pre>	<p>Enables confidential VMs or trusted launch for compute nodes. This option is not enabled by default.</p>	ConfidentialVM or TrustedLaunch .

Parameter	Description	Values
<pre>compute: platform: azure: settings: confidentialVM: uefiSettings: secureBoot:</pre>	Enables secure boot on compute nodes if you are using confidential VMs.	Enabled or Disabled . The default is Disabled .
<pre>compute: platform: azure: settings: confidentialVM: uefiSettings: virtualizedTrustedPlatformModule:</pre>	Enables the virtualized Trusted Platform Module (vTPM) feature on compute nodes if you are using confidential VMs.	Enabled or Disabled . The default is Disabled .
<pre>compute: platform: azure: settings: trustedLaunch: uefiSettings: secureBoot:</pre>	Enables secure boot on compute nodes if you are using trusted launch.	Enabled or Disabled . The default is Disabled .
<pre>compute: platform: azure: settings: trustedLaunch: uefiSettings: virtualizedTrustedPlatformModule:</pre>	Enables the vTPM feature on compute nodes if you are using trusted launch.	Enabled or Disabled . The default is Disabled .

Parameter	Description	Values
<pre>compute: platform: azure: osDisk: securityProfile: securityEncryptionT ype:</pre>	Enables the encryption of the virtual machine guest state for compute nodes. This parameter can only be used if you use Confidential VMs.	VMGuestStateOnly is the only supported value.
<pre>controlPlane: platform: azure: settings: securityType:</pre>	Enables confidential VMs or trusted launch for control plane nodes. This option is not enabled by default.	ConfidentialVM or TrustedLaunch .
<pre>controlPlane: platform: azure: settings: confidentialVM: uefiSettings: secureBoot:</pre>	Enables secure boot on control plane nodes if you are using confidential VMs.	Enabled or Disabled . The default is Disabled .
<pre>controlPlane: platform: azure: settings: confidentialVM: uefiSettings: virtualizedTrustedPI atformModule:</pre>	Enables the vTPM feature on control plane nodes if you are using confidential VMs.	Enabled or Disabled . The default is Disabled .

Parameter	Description	Values
<pre>controlPlane: platform: azure: settings: trustedLaunch: uefiSettings: secureBoot:</pre>	Enables secure boot on control plane nodes if you are using trusted launch.	Enabled or Disabled . The default is Disabled .
<pre>controlPlane: platform: azure: settings: trustedLaunch: uefiSettings: virtualizedTrustedPlatformModule:</pre>	Enables the vTPM feature on control plane nodes if you are using trusted launch.	Enabled or Disabled . The default is Disabled .
<pre>controlPlane: platform: azure: osDisk: securityProfile: securityEncryptionType:</pre>	Enables the encryption of the virtual machine guest state for control plane nodes. This parameter can only be used if you use Confidential VMs.	VMGuestStateOnly is the only supported value.
<pre>controlPlane: platform: azure: type:</pre>	Defines the Azure instance type for control plane machines.	String
<pre>controlPlane: platform: azure: zones:</pre>	The availability zones where the installation program creates control plane machines.	String list

Parameter	Description	Values
<pre>platform: azure: defaultMachinePlatform: settings: securityType:</pre>	Enables confidential VMs or trusted launch for all nodes. This option is not enabled by default.	ConfidentialVM or TrustedLaunch .
<pre>platform: azure: defaultMachinePlatform: settings: confidentialVM: uefiSettings: secureBoot:</pre>	Enables secure boot on all nodes if you are using confidential VMs.	Enabled or Disabled . The default is Disabled .
<pre>platform: azure: defaultMachinePlatform: settings: confidentialVM: uefiSettings: virtualizedTrustedPlatformModule:</pre>	Enables the virtualized Trusted Platform Module (vTPM) feature on all nodes if you are using confidential VMs.	Enabled or Disabled . The default is Disabled .
<pre>platform: azure: defaultMachinePlatform: settings: trustedLaunch: uefiSettings: secureBoot:</pre>	Enables secure boot on all nodes if you are using trusted launch.	Enabled or Disabled . The default is Disabled .

Parameter	Description	Values
<pre>platform: azure: defaultMachinePlatform: settings: trustedLaunch: uefiSettings: virtualizedTrustedPlatformModule:</pre>	Enables the vTPM feature on all nodes if you are using trusted launch.	Enabled or Disabled . The default is Disabled .
<pre>platform: azure: defaultMachinePlatform: osDisk: securityProfile: securityEncryptionType:</pre>	Enables the encryption of the virtual machine guest state for all nodes. This parameter can only be used if you use Confidential VMs.	VMGuestStateOnly is the only supported value.
<pre>platform: azure: defaultMachinePlatform: encryptionAtHost:</pre>	Enables host-level encryption for compute machines. You can enable this encryption alongside user-managed server-side encryption. This feature encrypts temporary, ephemeral, cached, and un-managed disks on the VM host. This parameter is not a prerequisite for user-managed server-side encryption.	true or false . The default is false .
<pre>platform: azure: defaultMachinePlatform: osDisk: diskEncryptionSetName:</pre>	The name of the disk encryption set that contains the encryption key from the installation prerequisites.	String, for example, production_disk_encryption_set .

Parameter	Description	Values
<pre>platform: azure: defaultMachinePlatform: osDisk: diskEncryptionSet: resourceGroup:</pre>	<p>The name of the Azure resource group that contains the disk encryption set from the installation prerequisites. To avoid deleting your Azure encryption key when the cluster is destroyed, this resource group must be different from the resource group where you install the cluster. This value is necessary only if you intend to install the cluster with user-managed disk encryption.</p>	<p>String, for example, production_encryption_resource_group.</p>
<pre>platform: azure: defaultMachinePlatform: osDisk: diskEncryptionSet: subscriptionId:</pre>	<p>Defines the Azure subscription of the disk encryption set where the disk encryption set resides. This secondary disk encryption set is used to encrypt compute machines.</p>	<p>String, in the format 00000000-0000-0000-0000-000000000000.</p>
<pre>platform: azure: defaultMachinePlatform: osDisk: diskSizeGB:</pre>	<p>The Azure disk size for the VM.</p>	<p>Integer that represents the size of the disk in GB. The default is 128.</p>
<pre>platform: azure: defaultMachinePlatform: osDisk: diskType:</pre>	<p>Defines the type of disk.</p>	<p>premium_LRS or standardSSD_LRS. The default is premium_LRS.</p>

Parameter	Description	Values
platform: azure: defaultMachinePlatform: osImage: publisher:	<p>Optional. By default, the installation program downloads and installs the Red Hat Enterprise Linux CoreOS (RHCOS) image that is used to boot control plane and compute machines. You can override the default behavior by using a custom RHCOS image that is available from the Azure Marketplace. The installation program uses this image for both types of machines.</p>	<p>String. The name of the image publisher.</p>
platform: azure: defaultMachinePlatform: osImage: offer:	<p>The name of Azure Marketplace offer that is associated with the custom RHCOS image. If you use platform.azure.defaultMachinePlatform.osImage.publisher, this field is required.</p>	<p>String. The name of the image offer.</p>
platform: azure: defaultMachinePlatform: osImage: sku:	<p>An instance of the Azure Marketplace offer. If you use platform.azure.defaultMachinePlatform.osImage.publisher, this field is required.</p>	<p>String. The SKU of the image offer.</p>
platform: azure: defaultMachinePlatform: osImage: version:	<p>The version number of the image SKU. If you use platform.azure.defaultMachinePlatform.osImage.publisher, this field is required.</p>	<p>String. The version of the image to use.</p>
platform: azure: defaultMachinePlatform: type:	<p>The Azure instance type for control plane and compute machines.</p>	<p>The Azure instance type.</p>

Parameter	Description	Values
<pre>platform: azure: defaultMachinePlatform: zones:</pre>	The availability zones where the installation program creates compute and control plane machines.	String list.
<pre>controlPlane: platform: azure: encryptionAtHost:</pre>	Enables host-level encryption for control plane machines. You can enable this encryption alongside user-managed server-side encryption. This feature encrypts temporary, ephemeral, cached and un-managed disks on the VM host. This is not a prerequisite for user-managed server-side encryption.	true or false . The default is false .
<pre>controlPlane: platform: azure: osDisk: diskEncryptionSet: resourceGroup:</pre>	The name of the Azure resource group that contains the disk encryption set from the installation prerequisites. This resource group should be different from the resource group where you install the cluster to avoid deleting your Azure encryption key when the cluster is destroyed. This value is only necessary if you intend to install the cluster with user-managed disk encryption.	String, for example production_encryption_resource_group .
<pre>controlPlane: platform: azure: osDisk: diskEncryptionSet: name:</pre>	The name of the disk encryption set that contains the encryption key from the installation prerequisites.	String, for example production_disk_encryption_set .
<pre>controlPlane: platform: azure: osDisk: diskEncryptionSet: subscriptionId:</pre>	Defines the Azure subscription of the disk encryption set where the disk encryption set resides. This secondary disk encryption set is used to encrypt control plane machines.	String, in the format 00000000-0000-0000-0000-000000000000 .

Parameter	Description	Values
controlPlane: platform: azure: osDisk: diskSizeGB:	The Azure disk size for the VM.	Integer that represents the size of the disk in GB. The default is 1024 .
controlPlane: platform: azure: osDisk: diskType:	Defines the type of disk.	premium_LRS or standardSSD_LRS . The default is premium_LRS .
controlPlane: platform: azure: osImage: publisher:	Optional. By default, the installation program downloads and installs the Red Hat Enterprise Linux CoreOS (RHCOS) image that is used to boot control plane machines. You can override the default behavior by using a custom RHCOS image that is available from the Azure Marketplace. The installation program uses this image for control plane machines only.	String. The name of the image publisher.
controlPlane: platform: azure: osImage: offer:	The name of Azure Marketplace offer that is associated with the custom RHCOS image. If you use controlPlane.platform.azure.osImage.publisher , this field is required.	String. The name of the image offer.
controlPlane: platform: azure: osImage: sku:	An instance of the Azure Marketplace offer. If you use controlPlane.platform.azure.osImage.publisher , this field is required.	String. The SKU of the image offer.
controlPlane: platform: azure: osImage: version:	The version number of the image SKU. If you use controlPlane.platform.azure.osImage.publisher , this field is required.	String. The version of the image to use.

Parameter	Description	Values
controlPlane: platform: azure: ultraSSDCapability:	Enables the use of Azure ultra disks for persistent storage on control plane machines. This requires that your Azure region and zone have ultra disks available.	Enabled, Disabled. The default is Disabled .
controlPlane: platform: azure: vmNetworkingType:	Enables accelerated networking. Accelerated networking enables single root I/O virtualization (SR-IOV) to a VM, improving its networking performance. If instance type of control plane machines support Accelerated networking, by default, the installer enables Accelerated networking, otherwise the default networking type is Basic .	Accelerated or Basic .
platform: azure: baseDomainResourceGroupName:	The name of the resource group that contains the DNS zone for your base domain.	String, for example production_cluster .
platform: azure: resourceGroupName:	The name of an already existing resource group to install your cluster to. This resource group must be empty and only used for this specific cluster; the cluster components assume ownership of all resources in the resource group. If you limit the service principal scope of the installation program to this resource group, you must ensure all other resources used by the installation program in your environment have the necessary permissions, such as the public DNS zone and virtual network. Destroying the cluster by using the installation program deletes this resource group.	String, for example existing_resource_group .
platform: azure: outboundType:	The outbound routing strategy used to connect your cluster to the internet. If you are using user-defined routing, you must have pre-existing networking available where the outbound routing has already been configured prior to installing a cluster. The installation program is not responsible for configuring user-defined routing. If	LoadBalancer, UserDefinedRouting, or NatGateway. The default is LoadBalancer .


Parameter	Description	Values
	<p>you specify the NatGateway routing strategy, the installation program will only create one NAT gateway. If you specify the NatGateway routing strategy, your account must have the Microsoft.Network/natGateways/read and Microsoft.Network/natGateways/write permissions.</p> <div data-bbox="488 477 592 1496" style="background-color: black; color: white; padding: 5px; font-weight: bold; text-align: center;">IMPORTANT</div> <p>NatGateway is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.</p> <p>For more information about the support scope of Red Hat Technology Preview features, see Technology Preview Features Support Scope.</p>	

Parameter	Description	Values
platform: azure: region:	The name of the Azure region that hosts your cluster.	Any valid region name, such as centralus .
platform: azure: zone:	List of availability zones to place machines in. For high availability, specify at least two zones.	List of zones, for example ["1", "2", "3"] .
platform: azure: customerManaged Key: keyVault: name:	Specifies the name of the key vault that contains the encryption key that is used to encrypt Azure storage.	String.
platform: azure: customerManaged Key: keyVault: keyName:	Specifies the name of the user-managed encryption key that is used to encrypt Azure storage.	String.
platform: azure: customerManaged Key: keyVault: resourceGroup:	Specifies the name of the resource group that contains the key vault and managed identity.	String.

Parameter	Description	Values
<pre>platform: azure: customerManagedKey: keyVault: subscriptionId:</pre>	Specifies the subscription ID that is associated with the key vault.	String, in the format 00000000-0000-0000-0000-000000000000 .
<pre>platform: azure: customerManagedKey: userAssignedIdentityKey:</pre>	Specifies the name of the user-assigned managed identity that resides in the resource group with the key vault and has access to the user-managed key.	String.
<pre>platform: azure: defaultMachinePlatform: ultraSSDCapability:</pre>	Enables the use of Azure ultra disks for persistent storage on control plane and compute machines. This requires that your Azure region and zone have ultra disks available.	Enabled, Disabled. The default is Disabled .
<pre>platform: azure: networkResourceGroupName:</pre>	The name of the resource group that contains the existing VNet that you want to deploy your cluster to. This name cannot be the same as the platform.azure.baseDomainResourceGroupName .	String.
<pre>platform: azure: virtualNetwork:</pre>	The name of the existing VNet that you want to deploy your cluster to.	String.

Parameter	Description	Values
<pre>platform: azure: controlPlaneSubnet :</pre>	The name of the existing subnet in your VNet that you want to deploy your control plane machines to.	Valid CIDR, for example 10.0.0.0/16 .
<pre>platform: azure: computeSubnet:</pre>	The name of the existing subnet in your VNet that you want to deploy your compute machines to.	Valid CIDR, for example 10.0.0.0/16 .
<pre>platform: azure: cloudName:</pre>	The name of the Azure cloud environment that is used to configure the Azure SDK with the appropriate Azure API endpoints. If empty, the default value AzurePublicCloud is used.	Any valid cloud environment, such as AzurePublicCloud or AzureUSGovernmentCloud .
<pre>platform: azure: defaultMachinePlatform: vmNetworkingType :</pre>	Enables accelerated networking. Accelerated networking enables single root I/O virtualization (SR-IOV) to a VM, improving its networking performance.	Accelerated or Basic . If instance type of control plane and compute machines support Accelerated networking, by default, the installer enables Accelerated networking, otherwise the default networking type is Basic .
<pre>operatorPublishingStrategy: apiserver:</pre>	Determines whether the load balancers that service the API are public or private. Set this parameter to Internal to prevent the API server from being accessible outside of your VNet. Set this parameter to External to make the API server accessible outside of your VNet. If you set this parameter, you must set the publish parameter to Mixed .	External or Internal . The default value is External .

Parameter	Description	Values
operatorPublishingStrategy: ingress:	Determines whether the DNS resources that the cluster creates for ingress traffic are publicly visible. Set this parameter to Internal to prevent the ingress VIP from being publicly accessible. Set this parameter to External to make the ingress VIP publicly accessible. If you set this parameter, you must set the publish parameter to Mixed .	External or Internal . The default value is External .



NOTE

You cannot customize [Azure Availability Zones](#) or [Use tags to organize your Azure resources](#) with an Azure cluster.

CHAPTER 8. INSTALLING ON AZURE STACK HUB

8.1. PREPARING TO INSTALL ON AZURE STACK HUB

8.1.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You have installed Azure Stack Hub version 2008 or later.

8.1.2. Requirements for installing OpenShift Container Platform on Azure Stack Hub

Before installing OpenShift Container Platform on Microsoft Azure Stack Hub, you must configure an Azure account.

See [Configuring an Azure Stack Hub account](#) for details about account configuration, account limits, DNS zone configuration, required roles, and creating service principals.

8.1.3. Choosing a method to install OpenShift Container Platform on Azure Stack Hub

You can install OpenShift Container Platform on installer-provisioned or user-provisioned infrastructure. The default installation type uses installer-provisioned infrastructure, where the installation program provisions the underlying infrastructure for the cluster. You can also install OpenShift Container Platform on infrastructure that you provision. If you do not use infrastructure that the installation program provisions, you must manage and maintain the cluster resources yourself.

See [Installation process](#) for more information about installer-provisioned and user-provisioned installation processes.

8.1.3.1. Installing a cluster on installer-provisioned infrastructure

You can install a cluster on Azure Stack Hub infrastructure that is provisioned by the OpenShift Container Platform installation program, by using the following method:

- [Installing a cluster on Azure Stack Hub with an installer-provisioned infrastructure](#) You can install OpenShift Container Platform on Azure Stack Hub infrastructure that is provisioned by the OpenShift Container Platform installation program.

8.1.3.2. Installing a cluster on user-provisioned infrastructure

You can install a cluster on Azure Stack Hub infrastructure that you provision, by using the following method:

- [Installing a cluster on Azure Stack Hub using ARM templates](#) You can install OpenShift Container Platform on Azure Stack Hub by using infrastructure that you provide. You can use the provided Azure Resource Manager (ARM) templates to assist with an installation.

8.1.4. Next steps

- [Configuring an Azure Stack Hub account](#)

8.2. CONFIGURING AN AZURE STACK HUB ACCOUNT

Before you can install OpenShift Container Platform, you must configure a Microsoft Azure account.



IMPORTANT

All Azure resources that are available through public endpoints are subject to resource name restrictions, and you cannot create resources that use certain terms. For a list of terms that Azure restricts, see [Resolve reserved resource name errors](#) in the Azure documentation.

8.2.1. Azure Stack Hub account limits

The OpenShift Container Platform cluster uses a number of Microsoft Azure Stack Hub components, and the default [Quota types in Azure Stack Hub](#) affect your ability to install OpenShift Container Platform clusters.

The following table summarizes the Azure Stack Hub components whose limits can impact your ability to install and run OpenShift Container Platform clusters.

Component	Number of components required by default	Description
vCPU	56	<p>A default cluster requires 56 vCPUs, so you must increase the account limit.</p> <p>By default, each cluster creates the following instances:</p> <ul style="list-style-type: none"> • One bootstrap machine, which is removed after installation • Three control plane machines • Three compute machines <p>Because the bootstrap, control plane, and worker machines use Standard_DS4_v2 virtual machines, which use 8 vCPUs, a default cluster requires 56 vCPUs. The bootstrap node VM is used only during installation.</p> <p>To deploy more worker nodes, enable autoscaling, deploy large workloads, or use a different instance type, you must further increase the vCPU limit for your account to ensure that your cluster can deploy the machines that you require.</p>
VNet	1	Each default cluster requires one Virtual Network (VNet), which contains two subnets.
Network interfaces	7	Each default cluster requires seven network interfaces. If you create more machines or your deployed workloads create load balancers, your cluster uses more network interfaces.

Component	Number of components required by default	Description						
Network security groups	2	<p>Each cluster creates network security groups for each subnet in the VNet. The default cluster creates network security groups for the control plane and for the compute node subnets:</p> <table border="1"> <tr> <td>control plane</td> <td>Allows the control plane machines to be reached on port 6443 from anywhere</td> </tr> <tr> <td>node</td> <td>Allows worker nodes to be reached from the internet on ports 80 and 443</td> </tr> </table>	control plane	Allows the control plane machines to be reached on port 6443 from anywhere	node	Allows worker nodes to be reached from the internet on ports 80 and 443		
control plane	Allows the control plane machines to be reached on port 6443 from anywhere							
node	Allows worker nodes to be reached from the internet on ports 80 and 443							
Network load balancers	3	<p>Each cluster creates the following load balancers:</p> <table border="1"> <tr> <td>default</td> <td>Public IP address that load balances requests to ports 80 and 443 across worker machines</td> </tr> <tr> <td>internal</td> <td>Private IP address that load balances requests to ports 6443 and 22623 across control plane machines</td> </tr> <tr> <td>external</td> <td>Public IP address that load balances requests to port 6443 across control plane machines</td> </tr> </table> <p>If your applications create more Kubernetes LoadBalancer service objects, your cluster uses more load balancers.</p>	default	Public IP address that load balances requests to ports 80 and 443 across worker machines	internal	Private IP address that load balances requests to ports 6443 and 22623 across control plane machines	external	Public IP address that load balances requests to port 6443 across control plane machines
default	Public IP address that load balances requests to ports 80 and 443 across worker machines							
internal	Private IP address that load balances requests to ports 6443 and 22623 across control plane machines							
external	Public IP address that load balances requests to port 6443 across control plane machines							
Public IP addresses	2	The public load balancer uses a public IP address. The bootstrap machine also uses a public IP address so that you can SSH into the machine to troubleshoot issues during installation. The IP address for the bootstrap node is used only during installation.						
Private IP addresses	7	The internal load balancer, each of the three control plane machines, and each of the three worker machines each use a private IP address.						

Additional resources

- [Optimizing storage](#).

8.2.2. Configuring a DNS zone in Azure Stack Hub

To successfully install OpenShift Container Platform on Azure Stack Hub, you must create DNS records in an Azure Stack Hub DNS zone. The DNS zone must be authoritative for the domain. To delegate a registrar's DNS zone to Azure Stack Hub, see Microsoft's documentation for [Azure Stack Hub datacenter DNS integration](#).

8.2.3. Required Azure Stack Hub roles

Your Microsoft Azure Stack Hub account must have the following roles for the subscription that you use:

- **Owner**

To set roles on the Azure portal, see the [Manage access to resources in Azure Stack Hub with role-based access control](#) in the Microsoft documentation.

8.2.4. Creating a service principal

Because OpenShift Container Platform and its installation program create Microsoft Azure resources by using the Azure Resource Manager, you must create a service principal to represent it.

Prerequisites

- Install or update the [Azure CLI](#).
- Your Azure account has the required roles for the subscription that you use.

Procedure

1. Register your environment:

```
$ az cloud register -n AzureStackCloud --endpoint-resource-manager <endpoint> 1
```

- 1 Specify the Azure Resource Manager endpoint, ``https://management.<region>.<fqdn>/``.

See the [Microsoft documentation](#) for details.

2. Set the active environment:

```
$ az cloud set -n AzureStackCloud
```

3. Update your environment configuration to use the specific API version for Azure Stack Hub:

```
$ az cloud update --profile 2019-03-01-hybrid
```

4. Log in to the Azure CLI:

```
$ az login
```

If you are in a multitenant environment, you must also supply the tenant ID.

5. If your Azure account uses subscriptions, ensure that you are using the right subscription:

- a. View the list of available accounts and record the **tenantId** value for the subscription you want to use for your cluster:

```
$ az account list --refresh
```

Example output


```
[
  {
    "cloudName": AzureStackCloud",
    "id": "9bab1460-96d5-40b3-a78e-17b15e978a80",
    "isDefault": true,
    "name": "Subscription Name",
    "state": "Enabled",
    "tenantId": "6057c7e9-b3ae-489d-a54e-de3f6bf6a8ee",
    "user": {
      "name": "you@example.com",
      "type": "user"
    }
  }
]
```

- b. View your active account details and confirm that the **tenantId** value matches the subscription you want to use:

```
$ az account show
```

Example output

```
{
  "environmentName": AzureStackCloud",
  "id": "9bab1460-96d5-40b3-a78e-17b15e978a80",
  "isDefault": true,
  "name": "Subscription Name",
  "state": "Enabled",
  "tenantId": "6057c7e9-b3ae-489d-a54e-de3f6bf6a8ee", 1
  "user": {
    "name": "you@example.com",
    "type": "user"
  }
}
```

- 1 Ensure that the value of the **tenantId** parameter is the correct subscription ID.

- c. If you are not using the right subscription, change the active subscription:

```
$ az account set -s <subscription_id> 1
```

- 1 Specify the subscription ID.

- d. Verify the subscription ID update:

```
$ az account show
```

Example output

```
{
  "environmentName": AzureStackCloud",
```

```

    "id": "33212d16-bdf6-45cb-b038-f6565b61edda",
    "isDefault": true,
    "name": "Subscription Name",
    "state": "Enabled",
    "tenantId": "8049c7e9-c3de-762d-a54e-dc3f6be6a7ee",
    "user": {
      "name": "you@example.com",
      "type": "user"
    }
  }
}

```

6. Record the **tenantId** and **id** parameter values from the output. You need these values during the OpenShift Container Platform installation.
7. Create the service principal for your account:

```

$ az ad sp create-for-rbac --role Contributor --name <service_principal> \ 1
--scopes /subscriptions/<subscription_id> 2
--years <years> 3

```

- 1 Specify the service principal name.
- 2 Specify the subscription ID.
- 3 Specify the number of years. By default, a service principal expires in one year. By using the **--years** option you can extend the validity of your service principal.

Example output

```

Creating 'Contributor' role assignment under scope '/subscriptions/<subscription_id>'
The output includes credentials that you must protect. Be sure that you do not
include these credentials in your code or check the credentials into your source
control. For more information, see https://aka.ms/azadsp-cli
{
  "appId": "ac461d78-bf4b-4387-ad16-7e32e328aec6",
  "displayName": <service_principal>,
  "password": "00000000-0000-0000-0000-000000000000",
  "tenantId": "8049c7e9-c3de-762d-a54e-dc3f6be6a7ee"
}

```

8. Record the values of the **appId** and **password** parameters from the previous output. You need these values during OpenShift Container Platform installation.

Additional resources

- For more information about CCO modes, see [About the Cloud Credential Operator](#).

8.2.5. Next steps

- Install an OpenShift Container Platform cluster:
 - [Installing a cluster quickly on Azure Stack Hub](#) .

- Install an OpenShift Container Platform cluster on Azure Stack Hub with user-provisioned infrastructure by following [Installing a cluster on Azure Stack Hub using ARM templates](#) .

8.3. INSTALLING A CLUSTER ON AZURE STACK HUB WITH AN INSTALLER-PROVISIONED INFRASTRUCTURE

In OpenShift Container Platform version 4.16, you can install a cluster on Microsoft Azure Stack Hub with an installer-provisioned infrastructure. However, you must manually configure the **install-config.yaml** file to specify values that are specific to Azure Stack Hub.



NOTE

While you can select **azure** when using the installation program to deploy a cluster using installer-provisioned infrastructure, this option is only supported for the Azure Public Cloud.

8.3.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [configured an Azure Stack Hub account](#) to host the cluster.
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.
- You verified that you have approximately 16 GB of local disk space. Installing the cluster requires that you download the RHCOS virtual hard disk (VHD) cluster image and upload it to your Azure Stack Hub environment so that it is accessible during deployment. Decompressing the VHD files requires this amount of local disk space.

8.3.2. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.



IMPORTANT

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

8.3.3. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.



NOTE

You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.



NOTE

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the **./openshift-install gather** command.

**NOTE**

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```

**NOTE**

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

8.3.4. Uploading the RHCOS cluster image

You must download the RHCOS virtual hard disk (VHD) cluster image and upload it to your Azure Stack Hub environment so that it is accessible during deployment.

Prerequisites

- Configure an Azure account.

Procedure

1. Obtain the RHCOS VHD cluster image:

- a. Export the URL of the RHCOS VHD to an environment variable.

```
$ export COMPRESSED_VHD_URL=$(openshift-install coreos print-stream-json | jq -r
'.architectures.x86_64.artifacts.azurestack.formats."vhd.gz".disk.location')
```

- b. Download the compressed RHCOS VHD file locally.

```
$ curl -O -L ${COMPRESSED_VHD_URL}
```

2. Decompress the VHD file.



NOTE

The decompressed VHD file is approximately 16 GB, so be sure that your host system has 16 GB of free space available. The VHD file can be deleted once you have uploaded it.

3. Upload the local VHD to the Azure Stack Hub environment, making sure that the blob is publicly available. For example, you can upload the VHD to a blob using the **az** cli or the web portal.

8.3.5. Obtaining the installation program

Before you install OpenShift Container Platform, download the installation file on the host you are using for installation.

Prerequisites

- You have a computer that runs Linux or macOS, with 500 MB of local disk space.

Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select **Azure** as the cloud provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.



IMPORTANT

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.



IMPORTANT

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.

4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#). This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

8.3.6. Manually creating the installation configuration file

Installing the cluster requires that you manually create the installation configuration file.


Prerequisites

- You have an SSH public key on your local machine to provide to the installation program. The key will be used for SSH authentication onto your cluster nodes for debugging and disaster recovery.
- You have obtained the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create an installation directory to store your required installation assets in:

```
$ mkdir <installation_directory>
```



IMPORTANT

You must create a directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

2. Customize the sample **install-config.yaml** file template that is provided and save it in the **<installation_directory>**.




NOTE

You must name this configuration file **install-config.yaml**.

Make the following modifications:

- a. Specify the required installation parameters.
- b. Update the **platform.azure** section to specify the parameters that are specific to Azure Stack Hub.

- c. Optional: Update one or more of the default configuration parameters to customize the installation.
For more information about the parameters, see "Installation configuration parameters".
3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the next step of the installation process. You must back it up now.

Additional resources

- [Installation configuration parameters for Azure Stack Hub](#)

8.3.6.1. Sample customized install-config.yaml file for Azure Stack Hub

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.



IMPORTANT

This sample YAML file is provided for reference only. Use it as a resource to enter parameter values into the installation configuration file that you created manually.

```

apiVersion: v1
baseDomain: example.com 1
credentialsMode: Manual
controlPlane: 2 3
  name: master
  platform:
    azure:
      osDisk:
        diskSizeGB: 1024 4
        diskType: premium_LRS
      replicas: 3
compute: 5
- name: worker
  platform:
    azure:
      osDisk:
        diskSizeGB: 512 6
        diskType: premium_LRS
      replicas: 3
metadata:
  name: test-cluster 7 8
networking:
  clusterNetwork:
    - cidr: 10.128.0.0/14
      hostPrefix: 23
  machineNetwork:
    - cidr: 10.0.0.0/16
  networkType: OVNKubernetes 9

```



```

serviceNetwork:
- 172.30.0.0/16
platform:
azure:
  armEndpoint: azurestack_arm_endpoint 10 11
  baseDomainResourceGroupName: resource_group 12 13
  region: azure_stack_local_region 14 15
  resourceGroupName: existing_resource_group 16
  outboundType: Loadbalancer
  cloudName: AzureStackCloud 17
  clusterOSImage: https://vhdsa.blob.example.example.com/vhd/rhcos-410.84.202112040202-0-
  azurestack.x86_64.vhd 18 19
  pullSecret: '{"auths": ...}' 20 21
  fips: false 22
  sshKey: ssh-ed25519 AAAA... 23
  additionalTrustBundle: | 24
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----

```

1 7 10 12 14 17 18 20 Required.

2 5 If you do not provide these parameters and values, the installation program provides the default value.

3 The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Although both sections currently define a single machine pool, it is possible that future versions of OpenShift Container Platform will support defining multiple compute pools during installation. Only one control plane pool is used.

4 6 You can specify the size of the disk to use in GB. Minimum recommendation for control plane nodes is 1024 GB.

8 The name of the cluster.

9 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.

11 The Azure Resource Manager endpoint that your Azure Stack Hub operator provides.

13 The name of the resource group that contains the DNS zone for your base domain.

15 The name of your Azure Stack Hub local region.


16 The name of an existing resource group to install your cluster to. If undefined, a new resource group is created for the cluster.

19 The URL of a storage blob in the Azure Stack environment that contains an RHCOS VHD.

21 The pull secret required to authenticate your cluster.

22 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography

modules that are provided with RHCOS instead.



IMPORTANT

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 23** You can optionally provide the **sshKey** value that you use to access the machines in your cluster.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- 24** If the Azure Stack Hub environment is using an internal Certificate Authority (CA), adding the CA certificate is required.

8.3.7. Manually manage cloud credentials

The Cloud Credential Operator (CCO) only supports your cloud provider in manual mode. As a result, you must specify the identity and access management (IAM) secrets for your cloud provider.

Procedure

1. If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

2. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

3. Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included 1 \
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml 2 \
  --to=<path_to_directory_for_credentials_requests> 3
```

- 1** The **--included** parameter includes only the manifests that your specific cluster configuration requires.

- 2 Specify the location of the **install-config.yaml** file.
- 3 Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

This command creates a YAML file for each **CredentialsRequest** object.

Sample **CredentialsRequest** object

```
apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: AzureProviderSpec
    roleBindings:
      - role: Contributor
    ...
```

4. Create YAML files for secrets in the **openshift-install** manifests directory that you generated previously. The secrets must be stored using the namespace and secret name defined in the **spec.secretRef** for each **CredentialsRequest** object.

Sample **CredentialsRequest** object with secrets

```
apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: AzureProviderSpec
    roleBindings:
      - role: Contributor
    ...
  secretRef:
    name: <component_secret>
    namespace: <component_namespace>
  ...
```


Sample **Secret** object

```
apiVersion: v1
kind: Secret
metadata:
  name: <component_secret>
  namespace: <component_namespace>
```

```

data:
  azure_subscription_id: <base64_encoded_azure_subscription_id>
  azure_client_id: <base64_encoded_azure_client_id>
  azure_client_secret: <base64_encoded_azure_client_secret>
  azure_tenant_id: <base64_encoded_azure_tenant_id>
  azure_resource_prefix: <base64_encoded_azure_resource_prefix>
  azure_resourcegroup: <base64_encoded_azure_resourcegroup>
  azure_region: <base64_encoded_azure_region>

```



IMPORTANT

Before upgrading a cluster that uses manually maintained credentials, you must ensure that the CCO is in an upgradeable state.

Additional resources

- [Updating cloud provider resources with manually maintained credentials](#)

8.3.8. Configuring the cluster to use an internal CA

If the Azure Stack Hub environment is using an internal Certificate Authority (CA), update the **cluster-proxy-01-config.yaml** file to configure the cluster to use the internal CA.

Prerequisites

- Create the **install-config.yaml** file and specify the certificate trust bundle in **.pem** format.
- Create the cluster manifests.

Procedure

1. From the directory in which the installation program creates files, go to the **manifests** directory.
2. Add **user-ca-bundle** to the **spec.trustedCA.name** field.

Example cluster-proxy-01-config.yaml file

```

apiVersion: config.openshift.io/v1
kind: Proxy
metadata:
  creationTimestamp: null
  name: cluster
spec:
  trustedCA:
    name: user-ca-bundle
status: {}

```

3. Optional: Back up the **manifests/ cluster-proxy-01-config.yaml** file. The installation program consumes the **manifests/** directory when you deploy the cluster.

8.3.9. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.

Procedure

- Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

- 1** For **<installation_directory>**, specify the location of your customized **./install-config.yaml** file.
- 2** To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to **<installation_directory>/./openshift_install.log**.



IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

8.3.10. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.



IMPORTANT

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**. To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

-

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.



NOTE

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your PATH.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

8.3.11. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

8.3.12. Logging in to the cluster by using the web console

The **kubeadmin** user exists by default after an OpenShift Container Platform installation. You can log in to your cluster as the **kubeadmin** user by using the OpenShift Container Platform web console.


Prerequisites

- You have access to the installation host.
- You completed a cluster installation and all cluster Operators are available.

Procedure

1. Obtain the password for the **kubeadmin** user from the **kubeadmin-password** file on the installation host:

```
$ cat <installation_directory>/auth/kubeadmin-password
```


**NOTE**

Alternatively, you can obtain the **kubeadmin** password from the **<installation_directory>/openshift_install.log** log file on the installation host.

- List the OpenShift Container Platform web console route:

```
$ oc get routes -n openshift-console | grep 'console-openshift'
```

**NOTE**

Alternatively, you can obtain the OpenShift Container Platform route from the **<installation_directory>/openshift_install.log** log file on the installation host.

Example output

```
console    console-openshift-console.apps.<cluster_name>.<base_domain>    console
https reencrypt/Redirect None
```

- Navigate to the route detailed in the output of the preceding command in a web browser and log in as the **kubeadmin** user.

Additional resources

- [Accessing the web console](#)

8.3.13. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- [About remote health monitoring](#)

8.3.14. Next steps

- [Validating an installation.](#)
- [Customize your cluster.](#)
- If necessary, you can [opt out of remote health reporting](#) .
- If necessary, you can [remove cloud provider credentials](#).

8.4. INSTALLING A CLUSTER ON AZURE STACK HUB WITH NETWORK CUSTOMIZATIONS

In OpenShift Container Platform version 4.16, you can install a cluster with a customized network configuration on infrastructure that the installation program provisions on Azure Stack Hub. By customizing your network configuration, your cluster can coexist with existing IP address allocations in your environment and integrate with existing MTU and VXLAN configurations.



NOTE

While you can select **azure** when using the installation program to deploy a cluster using installer-provisioned infrastructure, this option is only supported for the Azure Public Cloud.

8.4.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [configured an Azure Stack Hub account](#) to host the cluster.
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.
- You verified that you have approximately 16 GB of local disk space. Installing the cluster requires that you download the RHCOS virtual hard disk (VHD) cluster image and upload it to your Azure Stack Hub environment so that it is accessible during deployment. Decompressing the VHD files requires this amount of local disk space.

8.4.2. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.



IMPORTANT

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.


8.4.3. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes

through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.



NOTE

You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.



NOTE

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.

**NOTE**

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```

**NOTE**

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

8.4.4. Uploading the RHCOS cluster image

You must download the RHCOS virtual hard disk (VHD) cluster image and upload it to your Azure Stack Hub environment so that it is accessible during deployment.

Prerequisites

- Configure an Azure account.

Procedure

1. Obtain the RHCOS VHD cluster image:
 - a. Export the URL of the RHCOS VHD to an environment variable.

```
$ export COMPRESSED_VHD_URL=$(openshift-install coreos print-stream-json | jq -r '.architectures.x86_64.artifacts.azurestack.formats."vhd.gz".disk.location')
```

- b. Download the compressed RHCOS VHD file locally.

```
$ curl -O -L ${COMPRESSED_VHD_URL}
```

2. Decompress the VHD file.



NOTE

The decompressed VHD file is approximately 16 GB, so be sure that your host system has 16 GB of free space available. The VHD file can be deleted once you have uploaded it.

3. Upload the local VHD to the Azure Stack Hub environment, making sure that the blob is publicly available. For example, you can upload the VHD to a blob using the **az** cli or the web portal.

8.4.5. Obtaining the installation program


Before you install OpenShift Container Platform, download the installation file on the host you are using for installation.

Prerequisites

- You have a computer that runs Linux or macOS, with 500 MB of local disk space.


Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select **Azure** as the cloud provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.



IMPORTANT

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.



IMPORTANT

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.

4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#) . This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

8.4.6. Manually creating the installation configuration file

Installing the cluster requires that you manually create the installation configuration file.


Prerequisites

- You have an SSH public key on your local machine to provide to the installation program. The key will be used for SSH authentication onto your cluster nodes for debugging and disaster recovery.
- You have obtained the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create an installation directory to store your required installation assets in:

```
$ mkdir <installation_directory>
```



IMPORTANT

You must create a directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

2. Customize the sample **install-config.yaml** file template that is provided and save it in the **<installation_directory>**.



NOTE

You must name this configuration file **install-config.yaml**.

Make the following modifications:

- a. Specify the required installation parameters.
 - b. Update the **platform.azure** section to specify the parameters that are specific to Azure Stack Hub.
 - c. Optional: Update one or more of the default configuration parameters to customize the installation.
For more information about the parameters, see "Installation configuration parameters".
3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.

**IMPORTANT**


The **install-config.yaml** file is consumed during the next step of the installation process. You must back it up now.

Additional resources

- [Installation configuration parameters for Azure Stack Hub](#)

8.4.6.1. Sample customized install-config.yaml file for Azure Stack Hub

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.

**IMPORTANT**

This sample YAML file is provided for reference only. Use it as a resource to enter parameter values into the installation configuration file that you created manually.

```

apiVersion: v1
baseDomain: example.com 1
credentialsMode: Manual
controlPlane: 2 3
  name: master
  platform:
    azure:
      osDisk:
        diskSizeGB: 1024 4
        diskType: premium_LRS
      replicas: 3
compute: 5
- name: worker
  platform:
    azure:
      osDisk:
        diskSizeGB: 512 6
        diskType: premium_LRS
      replicas: 3
metadata:
  name: test-cluster 7 8
networking:
  clusterNetwork:
    - cidr: 10.128.0.0/14
      hostPrefix: 23
  machineNetwork:
    - cidr: 10.0.0.0/16
  networkType: OVNKubernetes 9
  serviceNetwork:
    - 172.30.0.0/16
platform:
  azure:
    armEndpoint: azurestack_arm_endpoint 10 11
    baseDomainResourceGroupName: resource_group 12 13

```

```

region: azure_stack_local_region 14 15
resourceGroupName: existing_resource_group 16
outboundType: Loadbalancer
cloudName: AzureStackCloud 17
clusterOSimage: https://vhdsa.blob.example.example.com/vhd/rhcos-410.84.202112040202-0-
azurestack.x86_64.vhd 18 19
pullSecret: '{"auths": ...}' 20 21
fips: false 22
sshKey: ssh-ed25519 AAAA... 23
additionalTrustBundle: | 24
  -----BEGIN CERTIFICATE-----
  <MY_TRUSTED_CA_CERT>
  -----END CERTIFICATE-----

```

1 7 10 12 14 17 18 20 Required.

2 5 If you do not provide these parameters and values, the installation program provides the default value.

3 The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Although both sections currently define a single machine pool, it is possible that future versions of OpenShift Container Platform will support defining multiple compute pools during installation. Only one control plane pool is used.

4 6 You can specify the size of the disk to use in GB. Minimum recommendation for control plane nodes is 1024 GB.

8 The name of the cluster.

9 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.

11 The Azure Resource Manager endpoint that your Azure Stack Hub operator provides.

13 The name of the resource group that contains the DNS zone for your base domain.

15 The name of your Azure Stack Hub local region.

16 The name of an existing resource group to install your cluster to. If undefined, a new resource group is created for the cluster.

19 The URL of a storage blob in the Azure Stack environment that contains an RHCOS VHD.

21 The pull secret required to authenticate your cluster.

22 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.



IMPORTANT

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 23 You can optionally provide the **sshKey** value that you use to access the machines in your cluster.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- 24 If the Azure Stack Hub environment is using an internal Certificate Authority (CA), adding the CA certificate is required.

8.4.7. Manually manage cloud credentials

The Cloud Credential Operator (CCO) only supports your cloud provider in manual mode. As a result, you must specify the identity and access management (IAM) secrets for your cloud provider.

Procedure

1. If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

2. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

3. Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included 1 \
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml 2 \
  --to=<path_to_directory_for_credentials_requests> 3
```

- 1 The **--included** parameter includes only the manifests that your specific cluster configuration requires.

- 2 Specify the location of the **install-config.yaml** file.

- 3 Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

This command creates a YAML file for each **CredentialsRequest** object.

Sample **CredentialsRequest** object

```

apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: AzureProviderSpec
    roleBindings:
      - role: Contributor
    ...

```

4. Create YAML files for secrets in the **openshift-install** manifests directory that you generated previously. The secrets must be stored using the namespace and secret name defined in the **spec.secretRef** for each **CredentialsRequest** object.

Sample **CredentialsRequest** object with secrets

```

apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: AzureProviderSpec
    roleBindings:
      - role: Contributor
    ...
  secretRef:
    name: <component_secret>
    namespace: <component_namespace>
  ...

```

Sample **Secret** object

```

apiVersion: v1
kind: Secret
metadata:
  name: <component_secret>
  namespace: <component_namespace>
data:
  azure_subscription_id: <base64_encoded_azure_subscription_id>

```

```

azure_client_id: <base64_encoded_azure_client_id>
azure_client_secret: <base64_encoded_azure_client_secret>
azure_tenant_id: <base64_encoded_azure_tenant_id>
azure_resource_prefix: <base64_encoded_azure_resource_prefix>
azure_resourcegroup: <base64_encoded_azure_resourcegroup>
azure_region: <base64_encoded_azure_region>

```



IMPORTANT

Before upgrading a cluster that uses manually maintained credentials, you must ensure that the CCO is in an upgradeable state.

Additional resources

- [Updating a cluster using the web console](#)
- [Updating a cluster using the CLI](#)

8.4.8. Configuring the cluster to use an internal CA

If the Azure Stack Hub environment is using an internal Certificate Authority (CA), update the **cluster-proxy-01-config.yaml** file to configure the cluster to use the internal CA.

Prerequisites

- Create the **install-config.yaml** file and specify the certificate trust bundle in **.pem** format.
- Create the cluster manifests.

Procedure

1. From the directory in which the installation program creates files, go to the **manifests** directory.
2. Add **user-ca-bundle** to the **spec.trustedCA.name** field.

Example cluster-proxy-01-config.yaml file

```

apiVersion: config.openshift.io/v1
kind: Proxy
metadata:
  creationTimestamp: null
  name: cluster
spec:
  trustedCA:
    name: user-ca-bundle
status: {}

```

3. Optional: Back up the **manifests/ cluster-proxy-01-config.yaml** file. The installation program consumes the **manifests/** directory when you deploy the cluster.

8.4.9. Network configuration phases

There are two phases prior to OpenShift Container Platform installation where you can customize the network configuration.

Phase 1

You can customize the following network-related fields in the **install-config.yaml** file before you create the manifest files:

- **networking.networkType**
- **networking.clusterNetwork**
- **networking.serviceNetwork**
- **networking.machineNetwork**

For more information on these fields, refer to *Installation configuration parameters*.



NOTE

Set the **networking.machineNetwork** to match the CIDR that the preferred NIC resides in.



IMPORTANT

The CIDR range **172.17.0.0/16** is reserved by libVirt. You cannot use this range or any range that overlaps with this range for any networks in your cluster.

Phase 2

After creating the manifest files by running **openshift-install create manifests**, you can define a customized Cluster Network Operator manifest with only the fields you want to modify. You can use the manifest to specify advanced network configuration.

You cannot override the values specified in phase 1 in the **install-config.yaml** file during phase 2. However, you can further customize the network plugin during phase 2.

8.4.10. Specifying advanced network configuration

You can use advanced network configuration for your network plugin to integrate your cluster into your existing network environment. You can specify advanced network configuration only before you install the cluster.



IMPORTANT

Customizing your network configuration by modifying the OpenShift Container Platform manifest files created by the installation program is not supported. Applying a manifest file that you create, as in the following procedure, is supported.

Prerequisites

- You have created the **install-config.yaml** file and completed any modifications to it.

Procedure

1. Change to the directory that contains the installation program and create the manifests:

```
$ ./openshift-install create manifests --dir <installation_directory> 1
```

- 1 **<installation_directory>** specifies the name of the directory that contains the **install-config.yaml** file for your cluster.
2. Create a stub manifest file for the advanced network configuration that is named **cluster-network-03-config.yml** in the **<installation_directory>/manifests/** directory:

```
apiVersion: operator.openshift.io/v1
kind: Network
metadata:
  name: cluster
spec:
```

3. Specify the advanced network configuration for your cluster in the **cluster-network-03-config.yml** file, such as in the following example:

Enable IPsec for the OVN-Kubernetes network provider

```
apiVersion: operator.openshift.io/v1
kind: Network
metadata:
  name: cluster
spec:
  defaultNetwork:
    ovnKubernetesConfig:
      ipsecConfig:
        mode: Full
```

4. Optional: Back up the **manifests/cluster-network-03-config.yml** file. The installation program consumes the **manifests/** directory when you create the Ignition config files.

8.4.11. Cluster Network Operator configuration

The configuration for the cluster network is specified as part of the Cluster Network Operator (CNO) configuration and stored in a custom resource (CR) object that is named **cluster**. The CR specifies the fields for the **Network** API in the **operator.openshift.io** API group.

The CNO configuration inherits the following fields during cluster installation from the **Network** API in the **Network.config.openshift.io** API group:

clusterNetwork

IP address pools from which pod IP addresses are allocated.

serviceNetwork

IP address pool for services.

defaultNetwork.type

Cluster network plugin. **OVNKubernetes** is the only supported plugin during installation.

You can specify the cluster network plugin configuration for your cluster by setting the fields for the **defaultNetwork** object in the CNO object named **cluster**.

8.4.11.1. Cluster Network Operator configuration object

The fields for the Cluster Network Operator (CNO) are described in the following table:

Table 8.1. Cluster Network Operator configuration object


Field	Type	Description
metadata.name	string	The name of the CNO object. This name is always cluster .
spec.clusterNetwork	array	<p>A list specifying the blocks of IP addresses from which pod IP addresses are allocated and the subnet prefix length assigned to each individual node in the cluster. For example:</p> <pre>spec: clusterNetwork: - cidr: 10.128.0.0/19 hostPrefix: 23 - cidr: 10.128.32.0/19 hostPrefix: 23</pre>
spec.serviceNetwork	array	<p>A block of IP addresses for services. The OpenShift SDN and OVN-Kubernetes network plugins support only a single IP address block for the service network. For example:</p> <pre>spec: serviceNetwork: - 172.30.0.0/14</pre> <p>You can customize this field only in the install-config.yaml file before you create the manifests. The value is read-only in the manifest file.</p>
spec.defaultNetwork	object	Configures the network plugin for the cluster network.
spec.kubeProxy Config	object	The fields for this object specify the kube-proxy configuration. If you are using the OVN-Kubernetes cluster network plugin, the kube-proxy configuration has no effect.

defaultNetwork object configuration

The values for the **defaultNetwork** object are defined in the following table:

Table 8.2. defaultNetwork object

Field	Type	Description
-------	------	-------------

Field	Type	Description
type	string	<p>OVNKubernetes. The Red Hat OpenShift Networking network plugin is selected during installation. This value cannot be changed after cluster installation.</p> <div style="display: flex; align-items: flex-start;">  <p>NOTE</p> <p>OpenShift Container Platform uses the OVN-Kubernetes network plugin by default. OpenShift SDN is no longer available as an installation choice for new clusters.</p> </div>
ovnKubernetesConfig	object	This object is only valid for the OVN-Kubernetes network plugin.

Configuration for the OVN-Kubernetes network plugin

The following table describes the configuration fields for the OVN-Kubernetes network plugin:

Table 8.3. ovnKubernetesConfig object

Field	Type	Description
mtu	integer	<p>The maximum transmission unit (MTU) for the Geneve (Generic Network Virtualization Encapsulation) overlay network. This is detected automatically based on the MTU of the primary network interface. You do not normally need to override the detected MTU.</p> <p>If the auto-detected value is not what you expect it to be, confirm that the MTU on the primary network interface on your nodes is correct. You cannot use this option to change the MTU value of the primary network interface on the nodes.</p> <p>If your cluster requires different MTU values for different nodes, you must set this value to 100 less than the lowest MTU value in your cluster. For example, if some nodes in your cluster have an MTU of 9001, and some have an MTU of 1500, you must set this value to 1400.</p>
genevePort	integer	The port to use for all Geneve packets. The default value is 6081 . This value cannot be changed after cluster installation.
ipsecConfig	object	Specify a configuration object for customizing the IPsec configuration.
ipv4	object	Specifies a configuration object for IPv4 settings.


Field	Type	Description
ipv6	object	Specifies a configuration object for IPv6 settings.
policyAuditConfig	object	Specify a configuration object for customizing network policy audit logging. If unset, the defaults audit log settings are used.
gatewayConfig	object	Optional: Specify a configuration object for customizing how egress traffic is sent to the node gateway.
		 <p>NOTE</p> <p>While migrating egress traffic, you can expect some disruption to workloads and service traffic until the Cluster Network Operator (CNO) successfully rolls out the changes.</p>

Table 8.4. `ovnKubernetesConfig.ipv4` object

Field	Type	Description
internalTransitSwitchSubnet	string	<p>If your existing network infrastructure overlaps with the 100.88.0.0/16 IPv4 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. The subnet for the distributed transit switch that enables east-west traffic. This subnet cannot overlap with any other subnets used by OVN-Kubernetes or on the host itself. It must be large enough to accommodate one IP address per node in your cluster.</p> <p>The default value is 100.88.0.0/16.</p>
internalJoinSubnet	string	<p>If your existing network infrastructure overlaps with the 100.64.0.0/16 IPv4 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. You must ensure that the IP address range does not overlap with any other subnet used by your OpenShift Container Platform installation. The IP address range must be larger than the maximum number of nodes that can be added to the cluster. For example, if the clusterNetwork.cidr value is 10.128.0.0/14 and the clusterNetwork.hostPrefix value is /23, then the maximum number of nodes is $2^{(23-14)}=512$.</p> <p>The default value is 100.64.0.0/16.</p>

Table 8.5. `ovnKubernetesConfig.ipv6` object

Field	Type	Description
-------	------	-------------

Field	Type	Description
internalTransitSwitchSubnet	string	<p>If your existing network infrastructure overlaps with the fd97::/64 IPv6 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. The subnet for the distributed transit switch that enables east-west traffic. This subnet cannot overlap with any other subnets used by OVN-Kubernetes or on the host itself. It must be large enough to accommodate one IP address per node in your cluster.</p> <p>The default value is fd97::/64.</p>
internalJoinSubnet	string	<p>If your existing network infrastructure overlaps with the fd98::/64 IPv6 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. You must ensure that the IP address range does not overlap with any other subnet used by your OpenShift Container Platform installation. The IP address range must be larger than the maximum number of nodes that can be added to the cluster.</p> <p>The default value is fd98::/64.</p>

Table 8.6. `policyAuditConfig` object

Field	Type	Description
rateLimit	integer	The maximum number of messages to generate every second per node. The default value is 20 messages per second.
maxFileSize	integer	The maximum size for the audit log in bytes. The default value is 50000000 or 50 MB.
maxLogFiles	integer	The maximum number of log files that are retained.
destination	string	<p>One of the following additional audit log targets:</p> <p>libc The libc syslog() function of the journald process on the host.</p> <p>udp:<host>:<port> A syslog server. Replace <host>:<port> with the host and port of the syslog server.</p> <p>unix:<file> A Unix Domain Socket file specified by <file>.</p> <p>null Do not send the audit logs to any additional target.</p>
syslogFacility	string	The syslog facility, such as kern , as defined by RFC5424. The default value is local0 .

Table 8.7. gatewayConfig object

Field	Type	Description
routingViaHost	boolean	<p>Set this field to true to send egress traffic from pods to the host networking stack. For highly-specialized installations and applications that rely on manually configured routes in the kernel routing table, you might want to route egress traffic to the host networking stack. By default, egress traffic is processed in OVN to exit the cluster and is not affected by specialized routes in the kernel routing table. The default value is false.</p> <p>This field has an interaction with the Open vSwitch hardware offloading feature. If you set this field to true, you do not receive the performance benefits of the offloading because egress traffic is processed by the host networking stack.</p>
ipForwarding	object	<p>You can control IP forwarding for all traffic on OVN-Kubernetes managed interfaces by using the ipForwarding specification in the Network resource. Specify Restricted to only allow IP forwarding for Kubernetes related traffic. Specify Global to allow forwarding of all IP traffic. For new installations, the default is Restricted. For updates to OpenShift Container Platform 4.14 or later, the default is Global.</p>
ipv4	object	<p>Optional: Specify an object to configure the internal OVN-Kubernetes masquerade address for host to service traffic for IPv4 addresses.</p>
ipv6	object	<p>Optional: Specify an object to configure the internal OVN-Kubernetes masquerade address for host to service traffic for IPv6 addresses.</p>

Table 8.8. gatewayConfig.ipv4 object

Field	Type	Description
internalMasqueradeSubnet	string	<p>The masquerade IPv4 addresses that are used internally to enable host to service traffic. The host is configured with these IP addresses as well as the shared gateway bridge interface. The default value is 169.254.169.0/29.</p>

Table 8.9. gatewayConfig.ipv6 object


Field	Type	Description
internalMasqueradeSubnet	string	<p>The masquerade IPv6 addresses that are used internally to enable host to service traffic. The host is configured with these IP addresses as well as the shared gateway bridge interface. The default value is fd69::/125.</p>

Table 8.10. `ipsecConfig` object

Field	Type	Description
<code>mode</code>	<code>string</code>	<p>Specifies the behavior of the IPsec implementation. Must be one of the following values:</p> <ul style="list-style-type: none"> ● Disabled: IPsec is not enabled on cluster nodes. ● External: IPsec is enabled for network traffic with external hosts. ● Full: IPsec is enabled for pod traffic and network traffic with external hosts.

Example OVN-Kubernetes configuration with IPsec enabled

```
defaultNetwork:
  type: OVNKubernetes
  ovnKubernetesConfig:
    mtu: 1400
    genevePort: 6081
    ipsecConfig:
      mode: Full
```




IMPORTANT

Using OVNKubernetes can lead to a stack exhaustion problem on IBM Power®.

kubeProxyConfig object configuration (OpenShiftSDN container network interface only)

The values for the `kubeProxyConfig` object are defined in the following table:

Table 8.11. `kubeProxyConfig` object

Field	Type	Description
<code>iptablesSyncPeriod</code>	<code>string</code>	<p>The refresh period for <code>iptables</code> rules. The default value is 30s. Valid suffixes include s, m, and h and are described in the Go time package documentation.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p>NOTE</p> <p>Because of performance improvements introduced in OpenShift Container Platform 4.3 and greater, adjusting the <code>iptablesSyncPeriod</code> parameter is no longer necessary.</p> </div> </div>

Field	Type	Description
proxyArguments.iptables-min-sync-period	array	<p>The minimum duration before refreshing iptables rules. This field ensures that the refresh does not happen too frequently. Valid suffixes include s, m, and h and are described in the Go time package. The default value is:</p> <pre>kubeProxyConfig: proxyArguments: iptables-min-sync-period: - 0s</pre>

8.4.12. Configuring hybrid networking with OVN-Kubernetes

You can configure your cluster to use hybrid networking with the OVN-Kubernetes network plugin. This allows a hybrid cluster that supports different node networking configurations.

Prerequisites

- You defined **OVNKubernetes** for the **networking.networkType** parameter in the **install-config.yaml** file. See the installation documentation for configuring OpenShift Container Platform network customizations on your chosen cloud provider for more information.

Procedure

- Change to the directory that contains the installation program and create the manifests:

```
$ ./openshift-install create manifests --dir <installation_directory>
```

where:

<installation_directory>

Specifies the name of the directory that contains the **install-config.yaml** file for your cluster.

- Create a stub manifest file for the advanced network configuration that is named **cluster-network-03-config.yml** in the **<installation_directory>/manifests/** directory:

```
$ cat <<EOF > <installation_directory>/manifests/cluster-network-03-config.yml
apiVersion: operator.openshift.io/v1
kind: Network
metadata:
  name: cluster
spec:
  EOF
```

where:

<installation_directory>

Specifies the directory name that contains the **manifests/** directory for your cluster.

- Open the **cluster-network-03-config.yml** file in an editor and configure OVN-Kubernetes with hybrid networking, such as in the following example:

Specify a hybrid networking configuration

```
apiVersion: operator.openshift.io/v1
kind: Network
metadata:
  name: cluster
spec:
  defaultNetwork:
    ovnKubernetesConfig:
      hybridOverlayConfig:
        hybridClusterNetwork: 1
        - cidr: 10.132.0.0/14
        hostPrefix: 23
```

- Specify the CIDR configuration used for nodes on the additional overlay network. The **hybridClusterNetwork** CIDR cannot overlap with the **clusterNetwork** CIDR.

- Save the **cluster-network-03-config.yml** file and quit the text editor.
- Optional: Back up the **manifests/cluster-network-03-config.yml** file. The installation program deletes the **manifests/** directory when creating the cluster.

8.4.13. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.

Procedure

- Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

- 1** For **<installation_directory>**, specify the location of your customized **./install-config.yaml** file.
- 2** To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to **<installation_directory>/openshift_install.log**.



IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

8.4.14. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.



IMPORTANT

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**.
To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:


```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.



NOTE

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your PATH.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

8.4.15. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```


-

- 1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

8.4.16. Logging in to the cluster by using the web console

The **kubeadmin** user exists by default after an OpenShift Container Platform installation. You can log in to your cluster as the **kubeadmin** user by using the OpenShift Container Platform web console.


Prerequisites

- You have access to the installation host.
- You completed a cluster installation and all cluster Operators are available.

Procedure

1. Obtain the password for the **kubeadmin** user from the **kubeadmin-password** file on the installation host:

```
$ cat <installation_directory>/auth/kubeadmin-password
```




NOTE

Alternatively, you can obtain the **kubeadmin** password from the **<installation_directory>/openshift_install.log** log file on the installation host.

2. List the OpenShift Container Platform web console route:

```
$ oc get routes -n openshift-console | grep 'console-openshift'
```



NOTE

Alternatively, you can obtain the OpenShift Container Platform route from the **<installation_directory>/openshift_install.log** log file on the installation host.

Example output

```
console    console-openshift-console.apps.<cluster_name>.<base_domain>    console
https reencrypt/Redirect None
```

3. Navigate to the route detailed in the output of the preceding command in a web browser and log in as the **kubeadmin** user.

Additional resources

- [Accessing the web console.](#)

8.4.17. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- [About remote health monitoring](#)


8.4.18. Next steps

- [Validating an installation.](#)
- [Customize your cluster.](#)
- If necessary, you can [opt out of remote health reporting](#) .
- If necessary, you can [remove cloud provider credentials](#).

8.5. INSTALLING A CLUSTER ON AZURE STACK HUB USING ARM TEMPLATES

In OpenShift Container Platform version 4.16, you can install a cluster on Microsoft Azure Stack Hub by using infrastructure that you provide.

Several [Azure Resource Manager](#) (ARM) templates are provided to assist in completing these steps or to help model your own.



IMPORTANT

The steps for performing a user-provisioned infrastructure installation are provided as an example only. Installing a cluster with infrastructure you provide requires knowledge of the cloud provider and the installation process of OpenShift Container Platform. Several ARM templates are provided to assist in completing these steps or to help model your own. You are also free to create the required resources through other methods; the templates are just an example.

8.5.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.

- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [configured an Azure Stack Hub account](#) to host the cluster.
- You downloaded the Azure CLI and installed it on your computer. See [Install the Azure CLI](#) in the Azure documentation. The documentation below was tested using version **2.28.0** of the Azure CLI. Azure CLI commands might perform differently based on the version you use.
- If you use a firewall and plan to use the Telemetry service, you [configured the firewall to allow the sites](#) that your cluster requires access to.

**NOTE**


Be sure to also review this site list if you are configuring a proxy.

8.5.2. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.

**IMPORTANT**

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

8.5.3. Configuring your Azure Stack Hub project

Before you can install OpenShift Container Platform, you must configure an Azure project to host it.

**IMPORTANT**

All Azure Stack Hub resources that are available through public endpoints are subject to resource name restrictions, and you cannot create resources that use certain terms. For a list of terms that Azure Stack Hub restricts, see [Resolve reserved resource name errors](#) in the Azure documentation.

8.5.3.1. Azure Stack Hub account limits

The OpenShift Container Platform cluster uses a number of Microsoft Azure Stack Hub components, and the default [Quota types in Azure Stack Hub](#) affect your ability to install OpenShift Container Platform clusters.

The following table summarizes the Azure Stack Hub components whose limits can impact your ability to install and run OpenShift Container Platform clusters.

Component	Number of components required by default	Description				
vCPU	56	<p>A default cluster requires 56 vCPUs, so you must increase the account limit.</p> <p>By default, each cluster creates the following instances:</p> <ul style="list-style-type: none"> ● One bootstrap machine, which is removed after installation ● Three control plane machines ● Three compute machines <p>Because the bootstrap, control plane, and worker machines use Standard_DS4_v2 virtual machines, which use 8 vCPUs, a default cluster requires 56 vCPUs. The bootstrap node VM is used only during installation.</p> <p>To deploy more worker nodes, enable autoscaling, deploy large workloads, or use a different instance type, you must further increase the vCPU limit for your account to ensure that your cluster can deploy the machines that you require.</p>				
VNet	1	Each default cluster requires one Virtual Network (VNet), which contains two subnets.				
Network interfaces	7	Each default cluster requires seven network interfaces. If you create more machines or your deployed workloads create load balancers, your cluster uses more network interfaces.				
Network security groups	2	<p>Each cluster creates network security groups for each subnet in the VNet. The default cluster creates network security groups for the control plane and for the compute node subnets:</p> <table border="1"> <tbody> <tr> <td>control plane</td> <td>Allows the control plane machines to be reached on port 6443 from anywhere</td> </tr> <tr> <td>node</td> <td>Allows worker nodes to be reached from the internet on ports 80 and 443</td> </tr> </tbody> </table>	control plane	Allows the control plane machines to be reached on port 6443 from anywhere	node	Allows worker nodes to be reached from the internet on ports 80 and 443
control plane	Allows the control plane machines to be reached on port 6443 from anywhere					
node	Allows worker nodes to be reached from the internet on ports 80 and 443					

Component	Number of components required by default	Description						
Network load balancers	3	<p>Each cluster creates the following load balancers:</p> <table border="1"> <tbody> <tr> <td>default</td> <td>Public IP address that load balances requests to ports 80 and 443 across worker machines</td> </tr> <tr> <td>internal</td> <td>Private IP address that load balances requests to ports 6443 and 22623 across control plane machines</td> </tr> <tr> <td>external</td> <td>Public IP address that load balances requests to port 6443 across control plane machines</td> </tr> </tbody> </table> <p>If your applications create more Kubernetes LoadBalancer service objects, your cluster uses more load balancers.</p>	default	Public IP address that load balances requests to ports 80 and 443 across worker machines	internal	Private IP address that load balances requests to ports 6443 and 22623 across control plane machines	external	Public IP address that load balances requests to port 6443 across control plane machines
default	Public IP address that load balances requests to ports 80 and 443 across worker machines							
internal	Private IP address that load balances requests to ports 6443 and 22623 across control plane machines							
external	Public IP address that load balances requests to port 6443 across control plane machines							
Public IP addresses	2	The public load balancer uses a public IP address. The bootstrap machine also uses a public IP address so that you can SSH into the machine to troubleshoot issues during installation. The IP address for the bootstrap node is used only during installation.						
Private IP addresses	7	The internal load balancer, each of the three control plane machines, and each of the three worker machines each use a private IP address.						

Additional resources

- [Optimizing storage](#).

8.5.3.2. Configuring a DNS zone in Azure Stack Hub

To successfully install OpenShift Container Platform on Azure Stack Hub, you must create DNS records in an Azure Stack Hub DNS zone. The DNS zone must be authoritative for the domain. To delegate a registrar's DNS zone to Azure Stack Hub, see Microsoft's documentation for [Azure Stack Hub datacenter DNS integration](#).

You can view Azure's DNS solution by visiting this [example for creating DNS zones](#).

8.5.3.3. Certificate signing requests management

Because your cluster has limited access to automatic machine management when you use infrastructure that you provision, you must provide a mechanism for approving cluster certificate signing requests (CSRs) after installation. The **kube-controller-manager** only approves the kubelet client CSRs. The **machine-approver** cannot guarantee the validity of a serving certificate that is requested by using kubelet credentials because it cannot confirm that the correct machine issued the request. You must determine and implement a method of verifying the validity of the kubelet serving certificate requests and approving them.

8.5.3.4. Required Azure Stack Hub roles

Your Microsoft Azure Stack Hub account must have the following roles for the subscription that you use:

- **Owner**

To set roles on the Azure portal, see the [Manage access to resources in Azure Stack Hub with role-based access control](#) in the Microsoft documentation.

8.5.3.5. Creating a service principal

Because OpenShift Container Platform and its installation program create Microsoft Azure resources by using the Azure Resource Manager, you must create a service principal to represent it.

Prerequisites

- Install or update the [Azure CLI](#).
- Your Azure account has the required roles for the subscription that you use.

Procedure

1. Register your environment:

```
$ az cloud register -n AzureStackCloud --endpoint-resource-manager <endpoint> 1
```

- 1** Specify the Azure Resource Manager endpoint, ``https://management.<region>.<fqdn>/``.

See the [Microsoft documentation](#) for details.

2. Set the active environment:

```
$ az cloud set -n AzureStackCloud
```

3. Update your environment configuration to use the specific API version for Azure Stack Hub:

```
$ az cloud update --profile 2019-03-01-hybrid
```

4. Log in to the Azure CLI:

```
$ az login
```

If you are in a multitenant environment, you must also supply the tenant ID.

5. If your Azure account uses subscriptions, ensure that you are using the right subscription:

- a. View the list of available accounts and record the **tenantId** value for the subscription you want to use for your cluster:

```
$ az account list --refresh
```

Example output

```
[
  {
    "cloudName": AzureStackCloud",
    "id": "9bab1460-96d5-40b3-a78e-17b15e978a80",
    "isDefault": true,
    "name": "Subscription Name",
    "state": "Enabled",
    "tenantId": "6057c7e9-b3ae-489d-a54e-de3f6bf6a8ee",
    "user": {
      "name": "you@example.com",
      "type": "user"
    }
  }
]
```

- b. View your active account details and confirm that the **tenantId** value matches the subscription you want to use:

```
$ az account show
```

Example output

```
{
  "environmentName": AzureStackCloud",
  "id": "9bab1460-96d5-40b3-a78e-17b15e978a80",
  "isDefault": true,
  "name": "Subscription Name",
  "state": "Enabled",
  "tenantId": "6057c7e9-b3ae-489d-a54e-de3f6bf6a8ee", 1
  "user": {
    "name": "you@example.com",
    "type": "user"
  }
}
```

- 1 Ensure that the value of the **tenantId** parameter is the correct subscription ID.

- c. If you are not using the right subscription, change the active subscription:

```
$ az account set -s <subscription_id> 1
```

- 1 Specify the subscription ID.

- d. Verify the subscription ID update:

```
$ az account show
```

Example output

```
{
  "environmentName": AzureStackCloud",
```

```

    "id": "33212d16-bdf6-45cb-b038-f6565b61edda",
    "isDefault": true,
    "name": "Subscription Name",
    "state": "Enabled",
    "tenantId": "8049c7e9-c3de-762d-a54e-dc3f6be6a7ee",
    "user": {
      "name": "you@example.com",
      "type": "user"
    }
  }
}

```

- Record the **tenantId** and **id** parameter values from the output. You need these values during the OpenShift Container Platform installation.
- Create the service principal for your account:

```

$ az ad sp create-for-rbac --role Contributor --name <service_principal> \ 1
--scopes /subscriptions/<subscription_id> 2
--years <years> 3

```

- Specify the service principal name.
- Specify the subscription ID.
- Specify the number of years. By default, a service principal expires in one year. By using the **--years** option you can extend the validity of your service principal.

Example output

```

Creating 'Contributor' role assignment under scope '/subscriptions/<subscription_id>'
The output includes credentials that you must protect. Be sure that you do not
include these credentials in your code or check the credentials into your source
control. For more information, see https://aka.ms/azadsp-cli
{
  "appId": "ac461d78-bf4b-4387-ad16-7e32e328aec6",
  "displayName": <service_principal>,
  "password": "00000000-0000-0000-0000-000000000000",
  "tenantId": "8049c7e9-c3de-762d-a54e-dc3f6be6a7ee"
}

```

- Record the values of the **appId** and **password** parameters from the previous output. You need these values during OpenShift Container Platform installation.

Additional resources

- For more information about CCO modes, see [About the Cloud Credential Operator](#).

8.5.4. Obtaining the installation program


Before you install OpenShift Container Platform, download the installation file on the host you are using for installation.

Prerequisites

- You have a computer that runs Linux or macOS, with 500 MB of local disk space.


Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select **Azure** as the cloud provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.



IMPORTANT

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.



IMPORTANT

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.

4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#). This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

8.5.5. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.

**IMPORTANT**

Do not skip this procedure in production environments, where disaster recovery and debugging is required.

**NOTE**

You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.

**NOTE**

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.

**NOTE**

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```

**NOTE**

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

8.5.6. Creating the installation files for Azure Stack Hub

To install OpenShift Container Platform on Microsoft Azure Stack Hub using user-provisioned infrastructure, you must generate the files that the installation program needs to deploy your cluster and modify them so that the cluster creates only the machines that it will use. You manually create the **install-config.yaml** file, and then generate and customize the Kubernetes manifests and Ignition config files. You also have the option to first set up a separate **var** partition during the preparation phases of installation.

8.5.6.1. Manually creating the installation configuration file

Installing the cluster requires that you manually create the installation configuration file.


Prerequisites

- You have an SSH public key on your local machine to provide to the installation program. The key will be used for SSH authentication onto your cluster nodes for debugging and disaster recovery.
- You have obtained the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create an installation directory to store your required installation assets in:


```
$ mkdir <installation_directory>
```



IMPORTANT

You must create a directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

2. Customize the sample **install-config.yaml** file template that is provided and save it in the **<installation_directory>**.



NOTE

You must name this configuration file **install-config.yaml**.

Make the following modifications for Azure Stack Hub:

- a. Set the **replicas** parameter to **0** for the **compute** pool:

```
compute:
- hyperthreading: Enabled
  name: worker
  platform: {}
  replicas: 0 1
```

- 1** Set to **0**.

The compute machines will be provisioned manually later.

- b. Update the **platform.azure** section of the **install-config.yaml** file to configure your Azure Stack Hub configuration:

```
platform:
  azure:
    armEndpoint: <azurestack_arm_endpoint> 1
    baseDomainResourceGroupName: <resource_group> 2
    cloudName: AzureStackCloud 3
    region: <azurestack_region> 4
```

- 1** Specify the Azure Resource Manager endpoint of your Azure Stack Hub environment, like **https://management.local.azurestack.external**.
- 2** Specify the name of the resource group that contains the DNS zone for your base domain.
- 3** Specify the Azure Stack Hub environment, which is used to configure the Azure SDK with the appropriate Azure API endpoints.
- 4** Specify the name of your Azure Stack Hub region.

3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.

**IMPORTANT**


The **install-config.yaml** file is consumed during the next step of the installation process. You must back it up now.

Additional resources

- [Installation configuration parameters for Azure Stack Hub](#)

8.5.6.2. Sample customized install-config.yaml file for Azure Stack Hub

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.

**IMPORTANT**

This sample YAML file is provided for reference only. Use it as a resource to enter parameter values into the installation configuration file that you created manually.

```

apiVersion: v1
baseDomain: example.com
controlPlane: 1
  name: master
  platform:
    azure:
      osDisk:
        diskSizeGB: 1024 2
        diskType: premium_LRS
      replicas: 3
compute: 3
- name: worker
  platform:
    azure:
      osDisk:
        diskSizeGB: 512 4
        diskType: premium_LRS
      replicas: 0
metadata:
  name: test-cluster 5
networking:
  clusterNetwork:
    - cidr: 10.128.0.0/14
      hostPrefix: 23
  machineNetwork:
    - cidr: 10.0.0.0/16
  networkType: OVNKubernetes 6
  serviceNetwork:
    - 172.30.0.0/16
platform:
  azure:
    armEndpoint: azurestack_arm_endpoint 7
    baseDomainResourceGroupName: resource_group 8
    region: azure_stack_local_region 9

```

```

resourceGroupName: existing_resource_group 10
outboundType: Loadbalancer
cloudName: AzureStackCloud 11
pullSecret: '{"auths": ...}' 12
fips: false 13
additionalTrustBundle: | 14
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
sshKey: ssh-ed25519 AAAA... 15

```

- 1 3** The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.
- 2 4** You can specify the size of the disk to use in GB. Minimum recommendation for control plane nodes is 1024 GB.
- 5** Specify the name of the cluster.
- 6** The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 7** Specify the Azure Resource Manager endpoint that your Azure Stack Hub operator provides.
- 8** Specify the name of the resource group that contains the DNS zone for your base domain.
- 9** Specify the name of your Azure Stack Hub local region.
- 10** Specify the name of an already existing resource group to install your cluster to. If undefined, a new resource group is created for the cluster.
- 11** Specify the Azure Stack Hub environment as your target platform.
- 12** Specify the pull secret required to authenticate your cluster.
- 13** Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.



IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 14 If your Azure Stack Hub environment uses an internal certificate authority (CA), add the necessary certificate bundle in **.pem** format.
- 15 You can optionally provide the **sshKey** value that you use to access the machines in your cluster.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

8.5.6.3. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
additionalTrustBundle: | 4
  -----BEGIN CERTIFICATE-----
  <MY_TRUSTED_CA_CERT>
  -----END CERTIFICATE-----
additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.

**NOTE**

The installation program does not support the proxy **readinessEndpoints** field.

**NOTE**

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.

**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

8.5.6.4. Exporting common variables for ARM templates

You must export a common set of variables that are used with the provided Azure Resource Manager (ARM) templates used to assist in completing a user-provided infrastructure install on Microsoft Azure Stack Hub.



NOTE

Specific ARM templates can also require additional exported variables, which are detailed in their related procedures.

Prerequisites

- Obtain the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Export common variables found in the **install-config.yaml** to be used by the provided ARM templates:

```
$ export CLUSTER_NAME=<cluster_name> 1
$ export AZURE_REGION=<azure_region> 2
$ export SSH_KEY=<ssh_key> 3
$ export BASE_DOMAIN=<base_domain> 4
$ export BASE_DOMAIN_RESOURCE_GROUP=<base_domain_resource_group> 5
```

- 1 The value of the **.metadata.name** attribute from the **install-config.yaml** file.
- 2 The region to deploy the cluster into. This is the value of the **.platform.azure.region** attribute from the **install-config.yaml** file.
- 3 The SSH RSA public key file as a string. You must enclose the SSH key in quotes since it contains spaces. This is the value of the **.sshKey** attribute from the **install-config.yaml** file.
- 4 The base domain to deploy the cluster to. The base domain corresponds to the DNS zone that you created for your cluster. This is the value of the **.baseDomain** attribute from the **install-config.yaml** file.
- 5 The resource group where the DNS zone exists. This is the value of the **.platform.azure.baseDomainResourceGroupName** attribute from the **install-config.yaml** file.

For example:

```
$ export CLUSTER_NAME=test-cluster
$ export AZURE_REGION=centralus
$ export SSH_KEY="ssh-rsa xxx/xxx/xxx= user@email.com"
$ export BASE_DOMAIN=example.com
$ export BASE_DOMAIN_RESOURCE_GROUP=ocp-cluster
```

2. Export the kubeadmin credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

8.5.6.5. Creating the Kubernetes manifest and Ignition config files

Because you must modify some cluster definition files and manually start the cluster machines, you must generate the Kubernetes manifest and Ignition config files that the cluster needs to configure the machines.

The installation configuration file transforms into the Kubernetes manifests. The manifests wrap into the Ignition configuration files, which are later used to configure the cluster machines.



IMPORTANT

- The Ignition config files that the OpenShift Container Platform installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

Prerequisites

- You obtained the OpenShift Container Platform installation program.
- You created the **install-config.yaml** installation configuration file.

Procedure

1. Change to the directory that contains the OpenShift Container Platform installation program and generate the Kubernetes manifests for the cluster:

```
$ ./openshift-install create manifests --dir <installation_directory> 1
```

- 1 For **<installation_directory>**, specify the installation directory that contains the **install-config.yaml** file you created.

2. Remove the Kubernetes manifest files that define the control plane machines:

```
$ rm -f <installation_directory>/openshift/99_openshift-cluster-api_master-machines-*.yaml
```


By removing these files, you prevent the cluster from automatically generating control plane machines.

3. Remove the Kubernetes manifest files that define the control plane machine set:

```
$ rm -f <installation_directory>/openshift/99_openshift-machine-api_master-control-plane-machine-set.yaml
```

4. Remove the Kubernetes manifest files that define the worker machines:

```
$ rm -f <installation_directory>/openshift/99_openshift-cluster-api_worker-machineset-*.yaml
```



IMPORTANT

If you disabled the **MachineAPI** capability when installing a cluster on user-provisioned infrastructure, you must remove the Kubernetes manifest files that define the worker machines. Otherwise, your cluster fails to install.

Because you create and manage the worker machines yourself, you do not need to initialize these machines.

5. Check that the **mastersSchedulable** parameter in the `<installation_directory>/manifests/cluster-scheduler-02-config.yml` Kubernetes manifest file is set to **false**. This setting prevents pods from being scheduled on the control plane machines:
 - a. Open the `<installation_directory>/manifests/cluster-scheduler-02-config.yml` file.
 - b. Locate the **mastersSchedulable** parameter and ensure that it is set to **false**.
 - c. Save and exit the file.
6. Optional: If you do not want [the Ingress Operator](#) to create DNS records on your behalf, remove the **privateZone** and **publicZone** sections from the `<installation_directory>/manifests/cluster-dns-02-config.yml` DNS configuration file:

```
apiVersion: config.openshift.io/v1
kind: DNS
metadata:
  creationTimestamp: null
  name: cluster
spec:
  baseDomain: example.openshift.com
  privateZone: ❶
    id: mycluster-100419-private-zone
  publicZone: ❷
    id: example.openshift.com
status: {}
```

❶ ❷ Remove this section completely.

If you do so, you must add ingress DNS records manually in a later step.

7. Optional: If your Azure Stack Hub environment uses an internal certificate authority (CA), you must update the **.spec.trustedCA.name** field in the `<installation_directory>/manifests/cluster-proxy-01-config.yml` file to use **user-ca-bundle**:

```
...
spec:
  trustedCA:
    name: user-ca-bundle
...
```

Later, you must update your bootstrap ignition to include the CA.

8. When configuring Azure on user-provisioned infrastructure, you must export some common variables defined in the manifest files to use later in the Azure Resource Manager (ARM) templates:

- a. Export the infrastructure ID by using the following command:

```
$ export INFRA_ID=<infra_id> 1
```

- 1 The OpenShift Container Platform cluster has been assigned an identifier (**INFRA_ID**) in the form of **<cluster_name>-<random_string>**. This will be used as the base name for most resources created using the provided ARM templates. This is the value of the **.status.infrastructureName** attribute from the **manifests/cluster-infrastructure-02-config.yml** file.

- b. Export the resource group by using the following command:

```
$ export RESOURCE_GROUP=<resource_group> 1
```

- 1 All resources created in this Azure deployment exists as part of a [resource group](#). The resource group name is also based on the **INFRA_ID**, in the form of **<cluster_name>-<random_string>-rg**. This is the value of the **.status.platformStatus.azure.resourceGroupName** attribute from the **manifests/cluster-infrastructure-02-config.yml** file.

9. Manually create your cloud credentials.

- a. From the directory that contains the installation program, obtain details of the OpenShift Container Platform release image that your **openshift-install** binary is built to use:

```
$ openshift-install version
```

Example output

```
release image quay.io/openshift-release-dev/ocp-release:4.y.z-x86_64
```

- b. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

- c. Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included 1 \
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \
  --to=<path_to_directory_for_credentials_requests> 3
```

- 1 The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2 Specify the location of the **install-config.yaml** file.
- 3 Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

This command creates a YAML file for each **CredentialsRequest** object.

Sample **CredentialsRequest** object

```
apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  labels:
    controller-tools.k8s.io: "1.0"
  name: openshift-image-registry-azure
  namespace: openshift-cloud-credential-operator
spec:
  secretRef:
    name: installer-cloud-credentials
    namespace: openshift-image-registry
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: AzureProviderSpec
    roleBindings:
      - role: Contributor
```

- d. Create YAML files for secrets in the **openshift-install** manifests directory that you generated previously. The secrets must be stored using the namespace and secret name defined in the **spec.secretRef** for each **CredentialsRequest** object. The format for the secret data varies for each cloud provider.

Sample **secrets.yaml** file:

```
apiVersion: v1
kind: Secret
metadata:
  name: ${secret_name}
  namespace: ${secret_namespace}
stringData:
  azure_subscription_id: ${subscription_id}
  azure_client_id: ${app_id}
  azure_client_secret: ${client_secret}
  azure_tenant_id: ${tenant_id}
  azure_resource_prefix: ${cluster_name}
  azure_resourcegroup: ${resource_group}
  azure_region: ${azure_region}
```

- e. Create a **cco-configmap.yaml** file in the manifests directory with the Cloud Credential Operator (CCO) disabled:

Sample **ConfigMap** object

```

apiVersion: v1
kind: ConfigMap
metadata:
  name: cloud-credential-operator-config
  namespace: openshift-cloud-credential-operator
  annotations:
    release.openshift.io/create-only: "true"
data:
  disabled: "true"

```

- To create the Ignition configuration files, run the following command from the directory that contains the installation program:

```
$ ./openshift-install create ignition-configs --dir <installation_directory> 1
```

- For **<installation_directory>**, specify the same installation directory.

Ignition config files are created for the bootstrap, control plane, and compute nodes in the installation directory. The **kubeadmin-password** and **kubeconfig** files are created in the **./<installation_directory>/auth** directory:

```

.
├── auth
│   ├── kubeadmin-password
│   └── kubeconfig
├── bootstrap.ign
├── master.ign
├── metadata.json
└── worker.ign

```

Additional resources

- [Manually manage cloud credentials](#)

8.5.6.6. Optional: Creating a separate **/var** partition

It is recommended that disk partitioning for OpenShift Container Platform be left to the installer. However, there are cases where you might want to create separate partitions in a part of the filesystem that you expect to grow.


OpenShift Container Platform supports the addition of a single partition to attach storage to either the **/var** partition or a subdirectory of **/var**. For example:

- /var/lib/containers:** Holds container-related content that can grow as more images and containers are added to a system.
- /var/lib/etcd:** Holds data that you might want to keep separate for purposes such as performance optimization of etcd storage.
- /var:** Holds data that you might want to keep separate for purposes such as auditing.

Storing the contents of a **/var** directory separately makes it easier to grow storage for those areas as needed and reinstall OpenShift Container Platform at a later date and keep that data intact. With this method, you will not have to pull all your containers again, nor will you have to copy massive log files

when you update systems.

Because **/var** must be in place before a fresh installation of Red Hat Enterprise Linux CoreOS (RHCOS), the following procedure sets up the separate **/var** partition by creating a machine config manifest that is inserted during the **openshift-install** preparation phases of an OpenShift Container Platform installation.



IMPORTANT

If you follow the steps to create a separate **/var** partition in this procedure, it is not necessary to create the Kubernetes manifest and Ignition config files again as described later in this section.

Procedure

1. Create a directory to hold the OpenShift Container Platform installation files:

```
$ mkdir $HOME/clusterconfig
```

2. Run **openshift-install** to create a set of files in the **manifest** and **openshift** subdirectories. Answer the system questions as you are prompted:

```
$ openshift-install create manifests --dir $HOME/clusterconfig
```

Example output

```
? SSH Public Key ...
INFO Credentials loaded from the "myprofile" profile in file "/home/myuser/.aws/credentials"
INFO Consuming Install Config from target directory
INFO Manifests created in: $HOME/clusterconfig/manifests and
$HOME/clusterconfig/openshift
```

3. Optional: Confirm that the installation program created manifests in the **clusterconfig/openshift** directory:

```
$ ls $HOME/clusterconfig/openshift/
```

Example output

```
99_kubeadmin-password-secret.yaml
99_openshift-cluster-api_master-machines-0.yaml
99_openshift-cluster-api_master-machines-1.yaml
99_openshift-cluster-api_master-machines-2.yaml
...
```

4. Create a Butane config that configures the additional partition. For example, name the file **\$HOME/clusterconfig/98-var-partition.bu**, change the disk device name to the name of the storage device on the **worker** systems, and set the storage size as appropriate. This example places the **/var** directory on a separate partition:

```
variant: openshift
version: 4.16.0
metadata:
```

```

labels:
  machineconfiguration.openshift.io/role: worker
name: 98-var-partition
storage:
  disks:
    - device: /dev/disk/by-id/<device_name> ❶
      partitions:
        - label: var
          start_mib: <partition_start_offset> ❷
          size_mib: <partition_size> ❸
          number: 5
      filesystems:
        - device: /dev/disk/by-partlabel/var
          path: /var
          format: xfs
          mount_options: [defaults, prjquota] ❹
          with_mount_unit: true

```

- ❶ The storage device name of the disk that you want to partition.
- ❷ When adding a data partition to the boot disk, a minimum value of 25000 MiB (Mebibytes) is recommended. The root file system is automatically resized to fill all available space up to the specified offset. If no value is specified, or if the specified value is smaller than the recommended minimum, the resulting root file system will be too small, and future reinstalls of RHCOS might overwrite the beginning of the data partition.
- ❸ The size of the data partition in mebibytes.
- ❹ The **prjquota** mount option must be enabled for filesystems used for container storage.



NOTE

When creating a separate **/var** partition, you cannot use different instance types for worker nodes, if the different instance types do not have the same device name.

5. Create a manifest from the Butane config and save it to the **clusterconfig/openshift** directory. For example, run the following command:

```
$ butane $HOME/clusterconfig/98-var-partition.bu -o $HOME/clusterconfig/openshift/98-var-partition.yaml
```

6. Run **openshift-install** again to create Ignition configs from a set of files in the **manifest** and **openshift** subdirectories:

```
$ openshift-install create ignition-configs --dir $HOME/clusterconfig
$ ls $HOME/clusterconfig/
auth bootstrap.ign master.ign metadata.json worker.ign
```

Now you can use the Ignition config files as input to the installation procedures to install Red Hat Enterprise Linux CoreOS (RHCOS) systems.

8.5.7. Creating the Azure resource group

You must create a Microsoft Azure [resource group](#). This is used during the installation of your OpenShift Container Platform cluster on Azure Stack Hub.

Prerequisites

- Configure an Azure account.
- Generate the Ignition config files for your cluster.

Procedure

- Create the resource group in a supported Azure region:

```
$ az group create --name ${RESOURCE_GROUP} --location ${AZURE_REGION}
```

8.5.8. Uploading the RHCOS cluster image and bootstrap Ignition config file

The Azure client does not support deployments based on files existing locally. You must copy and store the RHCOS virtual hard disk (VHD) cluster image and bootstrap Ignition config file in a storage container so they are accessible during deployment.

Prerequisites

- Configure an Azure account.
- Generate the Ignition config files for your cluster.

Procedure

1. Create an Azure storage account to store the VHD cluster image:

```
$ az storage account create -g ${RESOURCE_GROUP} --location ${AZURE_REGION} --name ${CLUSTER_NAME}sa --kind Storage --sku Standard_LRS
```



WARNING


The Azure storage account name must be between 3 and 24 characters in length and use numbers and lower-case letters only. If your **CLUSTER_NAME** variable does not follow these restrictions, you must manually define the Azure storage account name. For more information on Azure storage account name restrictions, see [Resolve errors for storage account names](#) in the Azure documentation.

2. Export the storage account key as an environment variable:

```
$ export ACCOUNT_KEY=`az storage account keys list -g ${RESOURCE_GROUP} --account-name ${CLUSTER_NAME}sa --query "[0].value" -o tsv`
```

- Export the URL of the RHCOS VHD to an environment variable:

```
$ export COMPRESSED_VHD_URL=$(openshift-install coreos print-stream-json | jq -r
'.architectures.x86_64.artifacts.azurestack.formats."vhd.gz".disk.location')
```



IMPORTANT

The RHCOS images might not change with every release of OpenShift Container Platform. You must specify an image with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Use the image version that matches your OpenShift Container Platform version if it is available.

- Create the storage container for the VHD:

```
$ az storage container create --name vhd --account-name ${CLUSTER_NAME}sa --account-
key ${ACCOUNT_KEY}
```

- Download the compressed RHCOS VHD file locally:

```
$ curl -O -L ${COMPRESSED_VHD_URL}
```

- Decompress the VHD file.



NOTE

The decompressed VHD file is approximately 16 GB, so be sure that your host system has 16 GB of free space available. You can delete the VHD file after you upload it.

- Copy the local VHD to a blob:

```
$ az storage blob upload --account-name ${CLUSTER_NAME}sa --account-key
${ACCOUNT_KEY} -c vhd -n "rhcos.vhd" -f rhcos-<rhcos_version>-azurestack.x86_64.vhd
```

- Create a blob storage container and upload the generated **bootstrap.ign** file:

```
$ az storage container create --name files --account-name ${CLUSTER_NAME}sa --
account-key ${ACCOUNT_KEY}
```

```
$ az storage blob upload --account-name ${CLUSTER_NAME}sa --account-key
${ACCOUNT_KEY} -c "files" -f "<installation_directory>/bootstrap.ign" -n "bootstrap.ign"
```

8.5.9. Example for creating DNS zones

DNS records are required for clusters that use user-provisioned infrastructure. You should choose the DNS strategy that fits your scenario.

For this example, [Azure Stack Hub's datacenter DNS integration](#) is used, so you will create a DNS zone.

**NOTE**

The DNS zone is not required to exist in the same resource group as the cluster deployment and might already exist in your organization for the desired base domain. If that is the case, you can skip creating the DNS zone; be sure the installation config you generated earlier reflects that scenario.

Prerequisites

- Configure an Azure account.
- Generate the Ignition config files for your cluster.

Procedure

- Create the new DNS zone in the resource group exported in the **BASE_DOMAIN_RESOURCE_GROUP** environment variable:

```
$ az network dns zone create -g ${BASE_DOMAIN_RESOURCE_GROUP} -n
${CLUSTER_NAME}.${BASE_DOMAIN}
```

You can skip this step if you are using a DNS zone that already exists.

You can learn more about [configuring a DNS zone in Azure Stack Hub](#) by visiting that section.

8.5.10. Creating a VNet in Azure Stack Hub

You must create a virtual network (VNet) in Microsoft Azure Stack Hub for your OpenShift Container Platform cluster to use. You can customize the VNet to meet your requirements. One way to create the VNet is to modify the provided Azure Resource Manager (ARM) template.

**NOTE**

If you do not use the provided ARM template to create your Azure Stack Hub infrastructure, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- Configure an Azure account.
- Generate the Ignition config files for your cluster.

Procedure

1. Copy the template from the **ARM template for the VNet** section of this topic and save it as **01_vnet.json** in your cluster's installation directory. This template describes the VNet that your cluster requires.
2. Create the deployment by using the **az** CLI:

```
$ az deployment group create -g ${RESOURCE_GROUP} \
  --template-file "<installation_directory>/01_vnet.json" \
  --parameters baseName="${INFRA_ID}" 1
```

- 1** The base name to be used in resource names; this is usually the cluster's infrastructure ID.

8.5.10.1. ARM template for the VNet

You can use the following Azure Resource Manager (ARM) template to deploy the VNet that you need for your OpenShift Container Platform cluster:

Example 8.1. 01_vnet.json ARM template

link:https://raw.githubusercontent.com/openshift/installer/release-4.16/upi/azurestack/01_vnet.json

8.5.11. Deploying the RHCOS cluster image for the Azure Stack Hub infrastructure

You must use a valid Red Hat Enterprise Linux CoreOS (RHCOS) image for Microsoft Azure Stack Hub for your OpenShift Container Platform nodes.

Prerequisites

- Configure an Azure account.
- Generate the Ignition config files for your cluster.
- Store the RHCOS virtual hard disk (VHD) cluster image in an Azure storage container.
- Store the bootstrap Ignition config file in an Azure storage container.

Procedure

1. Copy the template from the **ARM template for image storage** section of this topic and save it as **02_storage.json** in your cluster's installation directory. This template describes the image storage that your cluster requires.
2. Export the RHCOS VHD blob URL as a variable:

```
$ export VHD_BLOB_URL=`az storage blob url --account-name ${CLUSTER_NAME}sa --
  account-key ${ACCOUNT_KEY} -c vhd -n "rhcos.vhd" -o tsv`
```

3. Deploy the cluster image:

```
$ az deployment group create -g ${RESOURCE_GROUP} \
  --template-file "<installation_directory>/02_storage.json" \
  --parameters vhdBlobURL="${VHD_BLOB_URL}" 1 \
  --parameters baseName="${INFRA_ID}" 2 \
  --parameters storageAccount="${CLUSTER_NAME}sa" 3 \
  --parameters architecture="<architecture>" 4
```

- 1 The blob URL of the RHCOS VHD to be used to create master and worker machines.
- 2 The base name to be used in resource names; this is usually the cluster's infrastructure ID.
- 3 The name of your Azure storage account.
- 4 Specify the system architecture. Valid values are **x64** (default) or **Arm64**.

8.5.11.1. ARM template for image storage

You can use the following Azure Resource Manager (ARM) template to deploy the stored Red Hat Enterprise Linux CoreOS (RHCOS) image that you need for your OpenShift Container Platform cluster:

Example 8.2. 02_storage.json ARM template

link:https://raw.githubusercontent.com/openshift/installer/release-4.16/upi/azurestack/02_storage.json

8.5.12. Networking requirements for user-provisioned infrastructure

All the Red Hat Enterprise Linux CoreOS (RHCOS) machines require networking to be configured in **inittamfs** during boot to fetch their Ignition config files.

8.5.12.1. Network connectivity requirements

You must configure the network connectivity between machines to allow OpenShift Container Platform cluster components to communicate. Each machine must be able to resolve the hostnames of all other machines in the cluster.

This section provides details about the ports that are required.



IMPORTANT

In connected OpenShift Container Platform environments, all nodes are required to have internet access to pull images for platform containers and provide telemetry data to Red Hat.

Table 8.12. Ports used for all-machine to all-machine communications

Protocol	Port	Description
ICMP	N/A	Network reachability tests
TCP	1936	Metrics
	9000-9999	Host level services, including the node exporter on ports 9100-9101 and the Cluster Version Operator on port 9099 .
	10250-10259	The default ports that Kubernetes reserves

Protocol	Port	Description
UDP	4789	VXLAN
	6081	Geneve
	9000-9999	Host level services, including the node exporter on ports 9100-9101 .
	500	IPsec IKE packets
	4500	IPsec NAT-T packets
	123	Network Time Protocol (NTP) on UDP port 123 If an external NTP time server is configured, you must open UDP port 123 .
TCP/UDP	30000-32767	Kubernetes node port
ESP	N/A	IPsec Encapsulating Security Payload (ESP)

Table 8.13. Ports used for all-machine to control plane communications

Protocol	Port	Description
TCP	6443	Kubernetes API

Table 8.14. Ports used for control plane machine to control plane machine communications

Protocol	Port	Description
TCP	2379-2380	etcd server and peer ports

8.5.13. Creating networking and load balancing components in Azure Stack Hub

You must configure networking and load balancing in Microsoft Azure Stack Hub for your OpenShift Container Platform cluster to use. One way to create these components is to modify the provided Azure Resource Manager (ARM) template.

Load balancing requires the following DNS records:

- An **api** DNS record for the API public load balancer in the DNS zone.
- An **api-int** DNS record for the API internal load balancer in the DNS zone.



NOTE

If you do not use the provided ARM template to create your Azure Stack Hub infrastructure, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- Configure an Azure account.
- Generate the Ignition config files for your cluster.
- Create and configure a VNet and associated subnets in Azure Stack Hub.

Procedure

1. Copy the template from the **ARM template for the network and load balancers** section of this topic and save it as **03_infra.json** in your cluster's installation directory. This template describes the networking and load balancing objects that your cluster requires.
2. Create the deployment by using the **az** CLI:

```
$ az deployment group create -g ${RESOURCE_GROUP} \
  --template-file "<installation_directory>/03_infra.json" \
  --parameters baseName="${INFRA_ID}" 1
```

- 1** The base name to be used in resource names; this is usually the cluster's infrastructure ID.

3. Create an **api** DNS record and an **api-int** DNS record. When creating the API DNS records, the **\${BASE_DOMAIN_RESOURCE_GROUP}** variable must point to the resource group where the DNS zone exists.

- a. Export the following variable:

```
$ export PUBLIC_IP=`az network public-ip list -g ${RESOURCE_GROUP} --query "[?
  name=='${INFRA_ID}-master-pip'] | [0].ipAddress" -o tsv`
```

- b. Export the following variable:

```
$ export PRIVATE_IP=`az network lb frontend-ip show -g "$RESOURCE_GROUP" --lb-
  name "${INFRA_ID}-internal" -n internal-lb-ip --query "privateIpAddress" -o tsv`
```

- c. Create the **api** DNS record in a new DNS zone:

```
$ az network dns record-set a add-record -g ${BASE_DOMAIN_RESOURCE_GROUP} -
  z ${CLUSTER_NAME}.${BASE_DOMAIN} -n api -a ${PUBLIC_IP} --ttl 60
```

If you are adding the cluster to an existing DNS zone, you can create the **api** DNS record in it instead:

```
$ az network dns record-set a add-record -g ${BASE_DOMAIN_RESOURCE_GROUP} -
  z ${BASE_DOMAIN} -n api.${CLUSTER_NAME} -a ${PUBLIC_IP} --ttl 60
```

d. Create the **api-int** DNS record in a new DNS zone:

```
$ az network dns record-set a add-record -g ${BASE_DOMAIN_RESOURCE_GROUP} -
z "${CLUSTER_NAME}.${BASE_DOMAIN}" -n api-int -a ${PRIVATE_IP} --ttl 60
```

If you are adding the cluster to an existing DNS zone, you can create the **api-int** DNS record in it instead:

```
$ az network dns record-set a add-record -g ${BASE_DOMAIN_RESOURCE_GROUP} -
z ${BASE_DOMAIN} -n api-int.${CLUSTER_NAME} -a ${PRIVATE_IP} --ttl 60
```

8.5.13.1. ARM template for the network and load balancers

You can use the following Azure Resource Manager (ARM) template to deploy the networking objects and load balancers that you need for your OpenShift Container Platform cluster:

Example 8.3. 03_infra.json ARM template

link:https://raw.githubusercontent.com/openshift/installer/release-4.16/upi/azurestack/03_infra.json

8.5.14. Creating the bootstrap machine in Azure Stack Hub

You must create the bootstrap machine in Microsoft Azure Stack Hub to use during OpenShift Container Platform cluster initialization. One way to create this machine is to modify the provided Azure Resource Manager (ARM) template.



NOTE

If you do not use the provided ARM template to create your bootstrap machine, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- Configure an Azure account.
- Generate the Ignition config files for your cluster.
- Create and configure a VNet and associated subnets in Azure Stack Hub.
- Create and configure networking and load balancers in Azure Stack Hub.
- Create control plane and compute roles.

Procedure

1. Copy the template from the **ARM template for the bootstrap machines** section of this topic and save it as **04_bootstrap.json** in your cluster's installation directory. This template describes the bootstrap machine that your cluster requires.

2. Export the bootstrap URL variable:

```
$ bootstrap_url_expiry=`date -u -d "10 hours" '+%Y-%m-%dT%H:%MZ`
```

```
$ export BOOTSTRAP_URL=`az storage blob generate-sas -c 'files' -n 'bootstrap.ign' --https-only --full-uri --permissions r --expiry $bootstrap_url_expiry --account-name ${CLUSTER_NAME}sa --account-key ${ACCOUNT_KEY} -o tsv`
```

3. Export the bootstrap ignition variable:

- a. If your environment uses a public certificate authority (CA), run this command:

```
$ export BOOTSTRAP_IGNITION=`jq -rcnM --arg v "3.2.0" --arg url ${BOOTSTRAP_URL} '{ignition:{version:$v,config:{replace:{source:$url}}}}' | base64 | tr -d '\n`
```

- b. If your environment uses an internal CA, you must add your PEM encoded bundle to the bootstrap ignition stub so that your bootstrap virtual machine can pull the bootstrap ignition from the storage account. Run the following commands, which assume your CA is in a file called **CA.pem**:

```
$ export CA="data:text/plain;charset=utf-8;base64,$(cat CA.pem |base64 |tr -d '\n')"
```

```
$ export BOOTSTRAP_IGNITION=`jq -rcnM --arg v "3.2.0" --arg url "$BOOTSTRAP_URL" --arg cert "$CA" '{ignition:{version:$v,security:{tls:{certificateAuthorities:[{source:$cert}]},config:{replace:{source:$url}}}}' | base64 | tr -d '\n`
```

4. Create the deployment by using the **az** CLI:

```
$ az deployment group create --verbose -g ${RESOURCE_GROUP} \
  --template-file "<installation_directory>/04_bootstrap.json" \
  --parameters bootstrapIgnition="${BOOTSTRAP_IGNITION}" \ 1
  --parameters baseName="${INFRA_ID}" \ 2
  --parameters diagnosticsStorageAccountName="${CLUSTER_NAME}sa" 3
```

- 1** The bootstrap Ignition content for the bootstrap cluster.
- 2** The base name to be used in resource names; this is usually the cluster's infrastructure ID.
- 3** The name of the storage account for your cluster.

8.5.14.1. ARM template for the bootstrap machine

You can use the following Azure Resource Manager (ARM) template to deploy the bootstrap machine that you need for your OpenShift Container Platform cluster:

Example 8.4. 04_bootstrap.json ARM template

```
link:https://raw.githubusercontent.com/openshift/installer/release-4.16/upi/azurestack/04\_bootstrap.json
```

8.5.15. Creating the control plane machines in Azure Stack Hub

You must create the control plane machines in Microsoft Azure Stack Hub for your cluster to use. One way to create these machines is to modify the provided Azure Resource Manager (ARM) template.

If you do not use the provided ARM template to create your control plane machines, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, consider contacting Red Hat support with your installation logs.

Prerequisites

- Configure an Azure account.
- Generate the Ignition config files for your cluster.
- Create and configure a VNet and associated subnets in Azure Stack Hub.
- Create and configure networking and load balancers in Azure Stack Hub.
- Create control plane and compute roles.
- Create the bootstrap machine.

Procedure

1. Copy the template from the **ARM template for control plane machines** section of this topic and save it as **05_masters.json** in your cluster's installation directory. This template describes the control plane machines that your cluster requires.
2. Export the following variable needed by the control plane machine deployment:

```
$ export MASTER_IGNITION=`cat <installation_directory>/master.ign | base64 | tr -d '\n'`
```

3. Create the deployment by using the **az** CLI:

```
$ az deployment group create -g ${RESOURCE_GROUP} \
  --template-file "<installation_directory>/05_masters.json" \
  --parameters masterIgnition="${MASTER_IGNITION}" \ 1
  --parameters baseName="${INFRA_ID}" \ 2
  --parameters diagnosticsStorageAccountName="${CLUSTER_NAME}sa" 3
```

- 1** The Ignition content for the control plane nodes (also known as the master nodes).
- 2** The base name to be used in resource names; this is usually the cluster's infrastructure ID.
- 3** The name of the storage account for your cluster.

8.5.15.1. ARM template for control plane machines

You can use the following Azure Resource Manager (ARM) template to deploy the control plane machines that you need for your OpenShift Container Platform cluster:

Example 8.5. 05_masters.json ARM template

link:https://raw.githubusercontent.com/openshift/installer/release-4.16/upi/azurestack/05_masters.json

8.5.16. Wait for bootstrap completion and remove bootstrap resources in Azure Stack Hub

After you create all of the required infrastructure in Microsoft Azure Stack Hub, wait for the bootstrap process to complete on the machines that you provisioned by using the Ignition config files that you generated with the installation program.

Prerequisites

- Configure an Azure account.
- Generate the Ignition config files for your cluster.
- Create and configure a VNet and associated subnets in Azure Stack Hub.
- Create and configure networking and load balancers in Azure Stack Hub.
- Create control plane and compute roles.
- Create the bootstrap machine.
- Create the control plane machines.

Procedure

1. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install wait-for bootstrap-complete --dir <installation_directory> \ 1
--log-level info 2
```

- 1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

- 2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

If the command exits without a **FATAL** warning, your production control plane has initialized.

2. Delete the bootstrap resources:

```
$ az network nsg rule delete -g ${RESOURCE_GROUP} --nsg-name ${INFRA_ID}-nsg --
name bootstrap_ssh_in
$ az vm stop -g ${RESOURCE_GROUP} --name ${INFRA_ID}-bootstrap
$ az vm deallocate -g ${RESOURCE_GROUP} --name ${INFRA_ID}-bootstrap
$ az vm delete -g ${RESOURCE_GROUP} --name ${INFRA_ID}-bootstrap --yes
$ az disk delete -g ${RESOURCE_GROUP} --name ${INFRA_ID}-bootstrap OSDisk --no-
wait --yes
$ az network nic delete -g ${RESOURCE_GROUP} --name ${INFRA_ID}-bootstrap-nic --no-
```

```
wait
$ az storage blob delete --account-key ${ACCOUNT_KEY} --account-name
${CLUSTER_NAME}sa --container-name files --name bootstrap.ign
$ az network public-ip delete -g ${RESOURCE_GROUP} --name ${INFRA_ID}-bootstrap-
ssh-pip
```



NOTE

If you do not delete the bootstrap server, installation may not succeed due to API traffic being routed to the bootstrap server.

8.5.17. Creating additional worker machines in Azure Stack Hub

You can create worker machines in Microsoft Azure Stack Hub for your cluster to use by launching individual instances discretely or by automated processes outside the cluster, such as auto scaling groups. You can also take advantage of the built-in cluster scaling mechanisms and the machine API in OpenShift Container Platform.

In this example, you manually launch one instance by using the Azure Resource Manager (ARM) template. Additional instances can be launched by including additional resources of type **06_workers.json** in the file.

If you do not use the provided ARM template to create your control plane machines, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, consider contacting Red Hat support with your installation logs.

Prerequisites

- Configure an Azure account.
- Generate the Ignition config files for your cluster.
- Create and configure a VNet and associated subnets in Azure Stack Hub.
- Create and configure networking and load balancers in Azure Stack Hub.
- Create control plane and compute roles.
- Create the bootstrap machine.
- Create the control plane machines.

Procedure

1. Copy the template from the **ARM template for worker machines** section of this topic and save it as **06_workers.json** in your cluster's installation directory. This template describes the worker machines that your cluster requires.
2. Export the following variable needed by the worker machine deployment:

```
$ export WORKER_IGNITION=`cat <installation_directory>/worker.ign | base64 | tr -d "\n"
```

3. Create the deployment by using the **az** CLI:

```
$ az deployment group create -g ${RESOURCE_GROUP} \
```

```
--template-file "<installation_directory>/06_workers.json" \
--parameters workerIgnition="{WORKER_IGNITION}" \ 1
--parameters baseName="{INFRA_ID}" 2
--parameters diagnosticsStorageAccountName="{CLUSTER_NAME}sa" 3
```

- 1 The Ignition content for the worker nodes.
- 2 The base name to be used in resource names; this is usually the cluster's infrastructure ID.
- 3 The name of the storage account for your cluster.

8.5.17.1. ARM template for worker machines

You can use the following Azure Resource Manager (ARM) template to deploy the worker machines that you need for your OpenShift Container Platform cluster:

Example 8.6. 06_workers.json ARM template

```
link:https://raw.githubusercontent.com/openshift/installer/release-4.16/upi/azurestack/06_workers.json[]
```

8.5.18. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.



IMPORTANT

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**. To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.



NOTE

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your **PATH**.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

8.5.19. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

8.5.20. Approving the certificate signing requests for your machines

When you add machines to a cluster, two pending certificate signing requests (CSRs) are generated for each machine that you added. You must confirm that these CSRs are approved or, if necessary, approve them yourself. The client requests must be approved first, followed by the server requests.

Prerequisites

- You added machines to your cluster.

Procedure

1. Confirm that the cluster recognizes the machines:

```
$ oc get nodes
```

Example output

```
NAME      STATUS    ROLES    AGE   VERSION
master-0  Ready    master   63m   v1.29.4
master-1  Ready    master   63m   v1.29.4
master-2  Ready    master   64m   v1.29.4
```

The output lists all of the machines that you created.



NOTE

The preceding output might not include the compute nodes, also known as worker nodes, until some CSRs are approved.

- Review the pending CSRs and ensure that you see the client requests with the **Pending** or **Approved** status for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

```
NAME      AGE   REQUESTOR                                     CONDITION
csr-8b2br  15m   system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper  Pending
csr-8vnps  15m   system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper  Pending
...
```

In this example, two machines are joining the cluster. You might see more approved CSRs in the list.

- If the CSRs were not approved, after all of the pending CSRs for the machines you added are in **Pending** status, approve the CSRs for your cluster machines:



NOTE

Because the CSRs rotate automatically, approve your CSRs within an hour of adding the machines to the cluster. If you do not approve them within an hour, the certificates will rotate, and more than two certificates will be present for each node. You must approve all of these certificates. After the client CSR is approved, the Kubelet creates a secondary CSR for the serving certificate, which requires manual approval. Then, subsequent serving certificate renewal requests are automatically approved by the **machine-approver** if the Kubelet requests a new certificate with identical parameters.



NOTE

For clusters running on platforms that are not machine API enabled, such as bare metal and other user-provisioned infrastructure, you must implement a method of automatically approving the kubelet serving certificate requests (CSRs). If a request is not approved, then the **oc exec**, **oc rsh**, and **oc logs** commands cannot succeed, because a serving certificate is required when the API server connects to the kubelet. Any operation that contacts the Kubelet endpoint requires this certificate approval to be in place. The method must watch for new CSRs, confirm that the CSR was submitted by the **node-bootstrap** service account in the **system:node** or **system:admin** groups, and confirm the identity of the node.

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{"\n"}\n{{end}}{{end}}' | xargs --no-run-if-empty oc adm certificate approve
```



NOTE

Some Operators might not become available until some CSRs are approved.

4. Now that your client requests are approved, you must review the server requests for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

```
NAME      AGE   REQUESTOR                                     CONDITION
csr-bfd72 5m26s system:node:ip-10-0-50-126.us-east-2.compute.internal
Pending
csr-c57lv 5m26s system:node:ip-10-0-95-157.us-east-2.compute.internal
Pending
...
```

5. If the remaining CSRs are not approved, and are in the **Pending** status, approve the CSRs for your cluster machines:

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{\n"}\n{{end}}\n{{end}}' | xargs oc adm certificate approve
```

6. After all client and server CSRs have been approved, the machines have the **Ready** status. Verify this by running the following command:

```
$ oc get nodes
```

Example output

```
NAME      STATUS  ROLES  AGE  VERSION
master-0  Ready   master 73m  v1.29.4
master-1  Ready   master 73m  v1.29.4
master-2  Ready   master 74m  v1.29.4
worker-0  Ready   worker 11m  v1.29.4
worker-1  Ready   worker 11m  v1.29.4
```



NOTE

It can take a few minutes after approval of the server CSRs for the machines to transition to the **Ready** status.

Additional information

- For more information on CSRs, see [Certificate Signing Requests](#).

8.5.21. Adding the Ingress DNS records

If you removed the DNS Zone configuration when creating Kubernetes manifests and generating Ignition configs, you must manually create DNS records that point at the Ingress load balancer. You can create either a wildcard ***.apps.{baseDomain}**, or specific records. You can use A, CNAME, and other records per your requirements.

Prerequisites

- You deployed an OpenShift Container Platform cluster on Microsoft Azure Stack Hub by using infrastructure that you provisioned.
- Install the OpenShift CLI (**oc**).
- Install or update the [Azure CLI](#).

Procedure

1. Confirm the Ingress router has created a load balancer and populated the **EXTERNAL-IP** field:

```
$ oc -n openshift-ingress get service router-default
```

Example output

NAME	TYPE	CLUSTER-IP	EXTERNAL-IP	PORT(S)	AGE
router-default	LoadBalancer	172.30.20.10	35.130.120.110	80:32288/TCP,443:31215/TCP	20

- Export the Ingress router IP as a variable:

```
$ export PUBLIC_IP_ROUTER=`oc -n openshift-ingress get service router-default --no-headers | awk '{print $4}'`
```

- Add a ***.apps** record to the DNS zone.

- If you are adding this cluster to a new DNS zone, run:

```
$ az network dns record-set a add-record -g ${BASE_DOMAIN_RESOURCE_GROUP} -z ${CLUSTER_NAME}.${BASE_DOMAIN} -n *.apps -a ${PUBLIC_IP_ROUTER} --ttl 300
```

- If you are adding this cluster to an already existing DNS zone, run:

```
$ az network dns record-set a add-record -g ${BASE_DOMAIN_RESOURCE_GROUP} -z ${BASE_DOMAIN} -n *.apps.${CLUSTER_NAME} -a ${PUBLIC_IP_ROUTER} --ttl 300
```

If you prefer to add explicit domains instead of using a wildcard, you can create entries for each of the cluster's current routes:

```
$ oc get --all-namespaces -o jsonpath='{range .items[*]}{range .status.ingress[*]}{.host}{"\n"}{end}{end}' routes
```

Example output

```
oauth-openshift.apps.cluster.basedomain.com
console-openshift-console.apps.cluster.basedomain.com
downloads-openshift-console.apps.cluster.basedomain.com
alertmanager-main-openshift-monitoring.apps.cluster.basedomain.com
prometheus-k8s-openshift-monitoring.apps.cluster.basedomain.com
```

8.5.22. Completing an Azure Stack Hub installation on user-provisioned infrastructure

After you start the OpenShift Container Platform installation on Microsoft Azure Stack Hub user-provisioned infrastructure, you can monitor the cluster events until the cluster is ready.

Prerequisites

- Deploy the bootstrap machine for an OpenShift Container Platform cluster on user-provisioned Azure Stack Hub infrastructure.
- Install the **oc** CLI and log in.

Procedure


- Complete the cluster installation:

```
$ ./openshift-install --dir <installation_directory> wait-for install-complete 1
```

Example output

```
INFO Waiting up to 30m0s for the cluster to initialize...
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service.

8.6. INSTALLATION CONFIGURATION PARAMETERS FOR AZURE STACK HUB

Before you deploy an OpenShift Container Platform cluster on Azure Stack Hub, you provide a customized **install-config.yaml** installation configuration file that describes the details for your environment.

8.6.1. Available installation configuration parameters for Azure Stack Hub

The following tables specify the required, optional, and Azure Stack Hub-specific installation configuration parameters that you can set as part of the installation process.



NOTE

After installation, you cannot modify these parameters in the **install-config.yaml** file.

8.6.1.1. Required configuration parameters

Required installation configuration parameters are described in the following table:

Table 8.15. Required parameters

Parameter	Description	Values
apiVersion:	The API version for the install-config.yaml content. The current version is v1 . The installation program may also support older API versions.	String
baseDomain:	The base domain of your cloud provider. The base domain is used to create routes to your OpenShift Container Platform cluster components. The full DNS name for your cluster is a combination of the baseDomain and metadata.name parameter values that uses the <metadata.name> . <baseDomain> format.	A fully-qualified domain or subdomain name, such as example.com .
metadata:	Kubernetes resource ObjectMeta , from which only the name parameter is consumed.	Object
metadata: name:	The name of the cluster. DNS records for the cluster are all subdomains of {{.metadata.name}} . {{.baseDomain}} .	String of lowercase letters, hyphens (-), and periods (.), such as dev .
platform:	The configuration for the specific platform upon which to perform the installation: aws , baremetal , azure , gcp , ibmcloud , nutanix , openstack , powervs , vsphere , or {} . For additional information about platform . <platform> parameters, consult the table for your specific platform that follows.	Object

Parameter	Description	Values
<code>pullSecret:</code>	Get a pull secret from Red Hat OpenShift Cluster Manager to authenticate downloading container images for OpenShift Container Platform components from services such as Quay.io.	<pre>{ "auths":{ "cloud.openshift.com":{ "auth":"b3Blb=", "email":"you@example.com" }, "quay.io":{ "auth":"b3Blb=", "email":"you@example.com" } } }</pre>

8.6.1.2. Network configuration parameters

You can customize your installation configuration based on the requirements of your existing network infrastructure. For example, you can expand the IP address block for the cluster network or provide different IP address blocks than the defaults.

Only IPv4 addresses are supported.




NOTE

Globalnet is not supported with Red Hat OpenShift Data Foundation disaster recovery solutions. For regional disaster recovery scenarios, ensure that you use a nonoverlapping range of private IP addresses for the cluster and service networks in each cluster.

Table 8.16. Network parameters

Parameter	Description	Values
<code>networking:</code>	The configuration for the cluster network.	Object <div style="display: flex; align-items: center; margin-top: 10px;"> <div> <p>NOTE</p> <p>You cannot modify parameters specified by the networking object after installation.</p> </div> </div>
<code>networking: networkType:</code>	The Red Hat OpenShift Networking network plugin to install.	OVNKubernetes. OVNKubernetes is a CNI plugin for Linux networks and hybrid networks that contain both Linux and Windows servers. The default value is OVNKubernetes .


Parameter	Description	Values
<code>networking: clusterNetwork:</code>	<p>The IP address blocks for pods.</p> <p>The default value is 10.128.0.0/14 with a host prefix of /23.</p> <p>If you specify multiple IP address blocks, the blocks must not overlap.</p>	<p>An array of objects. For example:</p> <pre>networking: clusterNetwork: - cidr: 10.128.0.0/14 hostPrefix: 23</pre>
<code>networking: clusterNetwork: cidr:</code>	<p>Required if you use networking.clusterNetwork. An IP address block.</p> <p>An IPv4 network.</p>	<p>An IP address block in Classless Inter-Domain Routing (CIDR) notation. The prefix length for an IPv4 block is between 0 and 32.</p>
<code>networking: clusterNetwork: hostPrefix:</code>	<p>The subnet prefix length to assign to each individual node. For example, if hostPrefix is set to 23 then each node is assigned a /23 subnet out of the given cidr. A hostPrefix value of 23 provides 510 ($2^{(32 - 23)} - 2$) pod IP addresses.</p>	<p>A subnet prefix.</p> <p>The default value is 23.</p>
<code>networking: serviceNetwork:</code>	<p>The IP address block for services. The default value is 172.30.0.0/16.</p> <p>The OVN-Kubernetes network plugins supports only a single IP address block for the service network.</p>	<p>An array with an IP address block in CIDR format. For example:</p> <pre>networking: serviceNetwork: - 172.30.0.0/16</pre>
<code>networking: machineNetwork:</code>	<p>The IP address blocks for machines.</p> <p>If you specify multiple IP address blocks, the blocks must not overlap.</p>	<p>An array of objects. For example:</p> <pre>networking: machineNetwork: - cidr: 10.0.0.0/16</pre>
<code>networking: machineNetwork: cidr:</code>	<p>Required if you use networking.machineNetwork. An IP address block. The default value is 10.0.0.0/16 for all platforms other than libvirt and IBM Power® Virtual Server. For libvirt, the default value is 192.168.126.0/24. For IBM Power® Virtual Server, the default value is 192.168.0.0/24.</p>	<p>An IP network block in CIDR notation.</p> <p>For example, 10.0.0.0/16.</p> <div style="display: flex; align-items: flex-start;">  <div> <p>NOTE</p> <p>Set the networking.machineNetwork to match the CIDR that the preferred NIC resides in.</p> </div> </div>


8.6.1.3. Optional configuration parameters

Optional installation configuration parameters are described in the following table:

Table 8.17. Optional parameters



Parameter	Description	Values
<code>additionalTrustBundle:</code>	A PEM-encoded X.509 certificate bundle that is added to the nodes' trusted certificate store. This trust bundle may also be used when a proxy has been configured.	String
<code>capabilities:</code>	Controls the installation of optional core cluster components. You can reduce the footprint of your OpenShift Container Platform cluster by disabling optional components. For more information, see the "Cluster capabilities" page in <i>Installing</i> .	String array
<code>capabilities: baselineCapabilitySet:</code>	Selects an initial set of optional capabilities to enable. Valid values are None , v4.11 , v4.12 and vCurrent . The default value is vCurrent .	String
<code>capabilities: additionalEnabledCapabilities:</code>	Extends the set of optional capabilities beyond what you specify in baselineCapabilitySet . You may specify multiple capabilities in this parameter.	String array
<code>cpuPartitioningMode:</code>	Enables workload partitioning, which isolates OpenShift Container Platform services, cluster management workloads, and infrastructure pods to run on a reserved set of CPUs. Workload partitioning can only be enabled during installation and cannot be disabled after installation. While this field enables workload partitioning, it does not configure workloads to use specific CPUs. For more information, see the <i>Workload partitioning</i> page in the <i>Scalability and Performance</i> section.	None or AllNodes . None is the default value.
<code>compute:</code>	The configuration for the machines that comprise the compute nodes.	Array of MachinePool objects.

Parameter	Description	Values
<code>compute: architecture:</code>	Determines the instruction set architecture of the machines in the pool. Currently, clusters with varied architectures are not supported. All pools must specify the same architecture. Valid values are amd64 (the default).	String
<code>compute: hyperthreading:</code>	Whether to enable or disable simultaneous multithreading, or hyperthreading , on compute machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.  IMPORTANT If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.	Enabled or Disabled
<code>compute: name:</code>	Required if you use compute . The name of the machine pool.	worker
<code>compute: platform:</code>	Required if you use compute . Use this parameter to specify the cloud provider to host the worker machines. This parameter value must match the controlPlane.platform parameter value.	aws, azure, gcp, ibmcloud, nutanix, openstack, powervs, vsphere, or {}
<code>compute: replicas:</code>	The number of compute machines, which are also known as worker machines, to provision.	A positive integer greater than or equal to 2 . The default value is 3 .

Parameter	Description	Values
<code>featureSet:</code>	Enables the cluster for a feature set. A feature set is a collection of OpenShift Container Platform features that are not enabled by default. For more information about enabling a feature set during installation, see "Enabling features using feature gates".	String. The name of the feature set to enable, such as TechPreviewNoUpgrade .
<code>controlPlane:</code>	The configuration for the machines that comprise the control plane.	Array of MachinePool objects.
<code>controlPlane: architecture:</code>	Determines the instruction set architecture of the machines in the pool. Currently, clusters with varied architectures are not supported. All pools must specify the same architecture. Valid values are amd64 (the default).	String
<code>controlPlane: hyperthreading:</code>	<p>Whether to enable or disable simultaneous multithreading, or hyperthreading, on control plane machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.</p> <div style="display: flex; align-items: center;">  <div> <p>IMPORTANT</p> <p>If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.</p> </div> </div>	Enabled or Disabled
<code>controlPlane: name:</code>	Required if you use controlPlane . The name of the machine pool.	master
<code>controlPlane: platform:</code>	Required if you use controlPlane . Use this parameter to specify the cloud provider that hosts the control plane machines. This parameter value must match the compute.platform parameter value.	aws, azure, gcp, ibmcloud, nutanix, openstack, powervs, vsphere , or {}

Parameter	Description	Values
controlPlane: replicas:	The number of control plane machines to provision.	Supported values are 3 , or 1 when deploying single-node OpenShift.
credentialsMode:	The Cloud Credential Operator (CCO) mode. If no mode is specified, the CCO dynamically tries to determine the capabilities of the provided credentials, with a preference for mint mode on the platforms where multiple modes are supported.	Mint, Passthrough, Manual or an empty string (""). ^[1]

Parameter	Description	Values
<p>fips:</p>	<p>Enable or disable FIPS mode. The default is false (disabled). If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.</p> <div data-bbox="488 589 592 1637" style="background-color: black; color: white; padding: 10px;"> <p>IMPORTANT</p> <p>To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see Installing the system in FIPS mode.</p> <p>When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.</p> </div> <div data-bbox="488 1686 592 1883" style="background-color: black; color: white; padding: 10px;"> <p>NOTE</p> <p>If you are using Azure File storage, you cannot enable FIPS mode.</p> </div>	<p>false or true</p>
<p>imageContentSources:</p>	<p>Sources and repositories for the release-image content.</p>	<p>Array of objects. Includes a source and, optionally, mirrors, as described in the following rows of this table.</p>

Parameter	Description	Values
<code>imageContentSources: source:</code>	Required if you use imageContentSources . Specify the repository that users refer to, for example, in image pull specifications.	String
<code>imageContentSources: mirrors:</code>	Specify one or more repositories that may also contain the same images.	Array of strings
<code>publish:</code>	How to publish or expose the user-facing endpoints of your cluster, such as the Kubernetes API, OpenShift routes.	<p>Internal or External. The default value is External.</p> <p>Setting this field to Internal is not supported on non-cloud platforms.</p> <div style="display: flex; align-items: flex-start;">  <div> <p>IMPORTANT</p> <p>If the value of the field is set to Internal, the cluster will become non-functional. For more information, refer to BZ#1953035.</p> </div> </div>
<code>sshKey:</code>	<p>The SSH key to authenticate access to your cluster machines.</p> <div style="display: flex; align-items: flex-start;">  <div> <p>NOTE</p> <p>For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your ssh-agent process uses.</p> </div> </div>	For example, sshKey: ssh-ed25519 AAAA...

1. Not all CCO modes are supported for all cloud providers. For more information about CCO modes, see the "Managing cloud provider credentials" entry in the *Authentication and authorization* content.

8.6.1.4. Additional Azure Stack Hub configuration parameters

Additional Azure configuration parameters are described in the following table:

Table 8.18. Additional Azure Stack Hub parameters

Parameter	Description	Values
<pre>compute: platform: azure: osDisk: diskSizeGB:</pre>	The Azure disk size for the VM.	Integer that represents the size of the disk in GB. The default is 128 .
<pre>compute: platform: azure: osDisk: diskType:</pre>	Defines the type of disk.	standard_LRS or premium_LRS . The default is premium_LRS .
<pre>compute: platform: azure: type:</pre>	Defines the azure instance type for compute machines.	String
<pre>controlPlane: platform: azure: osDisk: diskSizeGB:</pre>	The Azure disk size for the VM.	Integer that represents the size of the disk in GB. The default is 1024 .
<pre>controlPlane: platform: azure: osDisk: diskType:</pre>	Defines the type of disk.	premium_LRS .
<pre>controlPlane: platform: azure: type:</pre>	Defines the azure instance type for control plane machines.	String

Parameter	Description	Values
<pre>platform: azure: defaultMachinePlatform: osDisk: diskSizeGB:</pre>	The Azure disk size for the VM.	Integer that represents the size of the disk in GB. The default is 128 .
<pre>platform: azure: defaultMachinePlatform: osDisk: diskType:</pre>	Defines the type of disk.	standard_LRS or premium_LRS . The default is premium_LRS .
<pre>platform: azure: defaultMachinePlatform: type:</pre>	The Azure instance type for control plane and compute machines.	The Azure instance type.
<pre>platform: azure: armEndpoint:</pre>	The URL of the Azure Resource Manager endpoint that your Azure Stack Hub operator provides.	String
<pre>platform: azure: baseDomainResourceGroupName:</pre>	The name of the resource group that contains the DNS zone for your base domain.	String, for example production_cluster .
<pre>platform: azure: region:</pre>	The name of your Azure Stack Hub local region.	String

Parameter	Description	Values
platform: azure: resourceGroupName:	<p>The name of an already existing resource group to install your cluster to. This resource group must be empty and only used for this specific cluster; the cluster components assume ownership of all resources in the resource group. If you limit the service principal scope of the installation program to this resource group, you must ensure all other resources used by the installation program in your environment have the necessary permissions, such as the public DNS zone and virtual network. Destroying the cluster by using the installation program deletes this resource group.</p>	<p>String, for example existing_resource_group.</p>
platform: azure: outboundType:	<p>The outbound routing strategy used to connect your cluster to the internet. If you are using user-defined routing, you must have pre-existing networking available where the outbound routing has already been configured prior to installing a cluster. The installation program is not responsible for configuring user-defined routing.</p>	<p>LoadBalancer or UserDefinedRouting. The default is LoadBalancer.</p>
platform: azure: cloudName:	<p>The name of the Azure cloud environment that is used to configure the Azure SDK with the appropriate Azure API endpoints.</p>	<p>AzureStackCloud</p>
clusterOSImage:	<p>The URL of a storage blob in the Azure Stack environment that contains an RHCOS VHD.</p>	<p>String, for example, <code>https://vhdsa.blob.example.example.com/vhd/rhcos-410.84.202112040202-0-azurestack.x86_64.vhd</code></p>

8.7. UNINSTALLING A CLUSTER ON AZURE STACK HUB

You can remove a cluster that you deployed to Azure Stack Hub.

8.7.1. Removing a cluster that uses installer-provisioned infrastructure

You can remove a cluster that uses installer-provisioned infrastructure from your cloud.

**NOTE**

After uninstallation, check your cloud provider for any resources not removed properly, especially with User Provisioned Infrastructure (UPI) clusters. There might be resources that the installer did not create or that the installer is unable to access.

Prerequisites


- You have a copy of the installation program that you used to deploy the cluster.
- You have the files that the installation program generated when you created your cluster.

Procedure

1. From the directory that contains the installation program on the computer that you used to install the cluster, run the following command:

```
$. /openshift-install destroy cluster \  
--dir <installation_directory> --log-level info 1 2
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.
- 2** To view different details, specify **warn**, **debug**, or **error** instead of **info**.

**NOTE**

You must specify the directory that contains the cluster definition files for your cluster. The installation program requires the **metadata.json** file in this directory to delete the cluster.

2. Optional: Delete the **<installation_directory>** directory and the OpenShift Container Platform installation program.

CHAPTER 9. INSTALLING ON GCP

9.1. PREPARING TO INSTALL ON GCP

9.1.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).

9.1.2. Requirements for installing OpenShift Container Platform on GCP

Before installing OpenShift Container Platform on Google Cloud Platform (GCP), you must create a service account and configure a GCP project. See [Configuring a GCP project](#) for details about creating a project, enabling API services, configuring DNS, GCP account limits, and supported GCP regions.

If the cloud identity and access management (IAM) APIs are not accessible in your environment, or if you do not want to store an administrator-level credential secret in the **kube-system** namespace, see [Manually creating long-term credentials for GCP](#) for other options.

9.1.3. Choosing a method to install OpenShift Container Platform on GCP

You can install OpenShift Container Platform on installer-provisioned or user-provisioned infrastructure. The default installation type uses installer-provisioned infrastructure, where the installation program provisions the underlying infrastructure for the cluster. You can also install OpenShift Container Platform on infrastructure that you provision. If you do not use infrastructure that the installation program provisions, you must manage and maintain the cluster resources yourself.

See [Installation process](#) for more information about installer-provisioned and user-provisioned installation processes.

9.1.3.1. Installing a cluster on installer-provisioned infrastructure

You can install a cluster on GCP infrastructure that is provisioned by the OpenShift Container Platform installation program, by using one of the following methods:

- [Installing a cluster quickly on GCP](#) You can install OpenShift Container Platform on GCP infrastructure that is provisioned by the OpenShift Container Platform installation program. You can install a cluster quickly by using the default configuration options.
- [Installing a customized cluster on GCP](#) You can install a customized cluster on GCP infrastructure that the installation program provisions. The installation program allows for some customization to be applied at the installation stage. Many other customization options are available [post-installation](#).
- [Installing a cluster on GCP with network customizations](#) You can customize your OpenShift Container Platform network configuration during installation, so that your cluster can coexist with your existing IP address allocations and adhere to your network requirements.
- [Installing a cluster on GCP in a restricted network](#) You can install OpenShift Container Platform on GCP on installer-provisioned infrastructure by using an internal mirror of the installation release content. You can use this method to install a cluster that does not require an

active internet connection to obtain the software components. While you can install OpenShift Container Platform by using the mirrored content, your cluster still requires internet access to use the GCP APIs.

- **Installing a cluster into an existing Virtual Private Cloud** You can install OpenShift Container Platform on an existing GCP Virtual Private Cloud (VPC). You can use this installation method if you have constraints set by the guidelines of your company, such as limits on creating new accounts or infrastructure.
- **Installing a private cluster on an existing VPC** You can install a private cluster on an existing GCP VPC. You can use this method to deploy OpenShift Container Platform on an internal network that is not visible to the internet.

9.1.3.2. Installing a cluster on user-provisioned infrastructure

You can install a cluster on GCP infrastructure that you provision, by using one of the following methods:

- **Installing a cluster on GCP with user-provisioned infrastructure** You can install OpenShift Container Platform on GCP infrastructure that you provide. You can use the provided Deployment Manager templates to assist with the installation.
- **Installing a cluster with shared VPC on user-provisioned infrastructure in GCP** You can use the provided Deployment Manager templates to create GCP resources in a shared VPC infrastructure.
- **Installing a cluster on GCP in a restricted network with user-provisioned infrastructure** You can install OpenShift Container Platform on GCP in a restricted network with user-provisioned infrastructure. By creating an internal mirror of the installation release content, you can install a cluster that does not require an active internet connection to obtain the software components. You can also use this installation method to ensure that your clusters only use container images that satisfy your organizational controls on external content.

9.1.4. Next steps

- [Configuring a GCP project](#)

9.2. CONFIGURING A GCP PROJECT


Before you can install OpenShift Container Platform, you must configure a Google Cloud Platform (GCP) project to host it.

9.2.1. Creating a GCP project

To install OpenShift Container Platform, you must create a project in your Google Cloud Platform (GCP) account to host the cluster.

Procedure

- Create a project to host your OpenShift Container Platform cluster. See [Creating and Managing Projects](#) in the GCP documentation.



IMPORTANT

Your GCP project must use the Premium Network Service Tier if you are using installer-provisioned infrastructure. The Standard Network Service Tier is not supported for clusters installed using the installation program. The installation program configures internal load balancing for the **api-int.<cluster_name>.<base_domain>** URL; the Premium Tier is required for internal load balancing.

9.2.2. Enabling API services in GCP

Your Google Cloud Platform (GCP) project requires access to several API services to complete OpenShift Container Platform installation.

Prerequisites

- You created a project to host your cluster.

Procedure

- Enable the following required API services in the project that hosts your cluster. You may also enable optional API services which are not required for installation. See [Enabling services](#) in the GCP documentation.

Table 9.1. Required API services

API service	Console service name
Compute Engine API	compute.googleapis.com
Cloud Resource Manager API	cloudresourcemanager.googleapis.com
Google DNS API	dns.googleapis.com
IAM Service Account Credentials API	iamcredentials.googleapis.com
Identity and Access Management (IAM) API	iam.googleapis.com
Service Usage API	serviceusage.googleapis.com

Table 9.2. Optional API services

API service	Console service name
Google Cloud APIs	cloudapis.googleapis.com
Service Management API	servicemanagement.googleapis.com
Google Cloud Storage JSON API	storage-api.googleapis.com

API service	Console service name
Cloud Storage	storage-component.googleapis.com

9.2.3. Configuring DNS for GCP

To install OpenShift Container Platform, the Google Cloud Platform (GCP) account you use must have a dedicated public hosted zone in the same project that you host the OpenShift Container Platform cluster. This zone must be authoritative for the domain. The DNS service provides cluster DNS resolution and name lookup for external connections to the cluster.

Procedure

1. Identify your domain, or subdomain, and registrar. You can transfer an existing domain and registrar or obtain a new one through GCP or another source.



NOTE

If you purchase a new domain, it can take time for the relevant DNS changes to propagate. For more information about purchasing domains through Google, see [Google Domains](#).

2. Create a public hosted zone for your domain or subdomain in your GCP project. See [Creating public zones](#) in the GCP documentation. Use an appropriate root domain, such as **openshiftcorp.com**, or subdomain, such as **clusters.openshiftcorp.com**.
3. Extract the new authoritative name servers from the hosted zone records. See [Look up your Cloud DNS name servers](#) in the GCP documentation. You typically have four name servers.
4. Update the registrar records for the name servers that your domain uses. For example, if you registered your domain to Google Domains, see the following topic in the Google Domains Help: [How to switch to custom name servers](#).
5. If you migrated your root domain to Google Cloud DNS, migrate your DNS records. See [Migrating to Cloud DNS](#) in the GCP documentation.
6. If you use a subdomain, follow your company's procedures to add its delegation records to the parent domain. This process might include a request to your company's IT department or the division that controls the root domain and DNS services for your company.

9.2.4. GCP account limits

The OpenShift Container Platform cluster uses a number of Google Cloud Platform (GCP) components, but the default [Quotas](#) do not affect your ability to install a default OpenShift Container Platform cluster.

A default cluster, which contains three compute and three control plane machines, uses the following resources. Note that some resources are required only during the bootstrap process and are removed after the cluster deploys.

Table 9.3. GCP resources used in a default cluster

Service	Component	Location	Total resources required	Resources removed after bootstrap
Service account	IAM	Global	6	1
Firewall rules	Compute	Global	11	1
Forwarding rules	Compute	Global	2	0
In-use global IP addresses	Compute	Global	4	1
Health checks	Compute	Global	3	0
Images	Compute	Global	1	0
Networks	Compute	Global	2	0
Static IP addresses	Compute	Region	4	1
Routers	Compute	Global	1	0
Routes	Compute	Global	2	0
Subnetworks	Compute	Global	2	0
Target pools	Compute	Global	3	0
CPUs	Compute	Region	28	4
Persistent disk SSD (GB)	Compute	Region	896	128



NOTE

If any of the quotas are insufficient during installation, the installation program displays an error that states both which quota was exceeded and the region.

Be sure to consider your actual cluster size, planned cluster growth, and any usage from other clusters that are associated with your account. The CPU, static IP addresses, and persistent disk SSD (storage) quotas are the ones that are most likely to be insufficient.

If you plan to deploy your cluster in one of the following regions, you will exceed the maximum storage quota and are likely to exceed the CPU quota limit:

- **asia-east2**
- **asia-northeast2**

- **asia-south1**
- **australia-southeast1**
- **europa-north1**
- **europa-west2**
- **europa-west3**
- **europa-west6**
- **northamerica-northeast1**
- **southamerica-east1**
- **us-west2**

You can increase resource quotas from the [GCP console](#), but you might need to file a support ticket. Be sure to plan your cluster size early so that you can allow time to resolve the support ticket before you install your OpenShift Container Platform cluster.

9.2.5. Creating a service account in GCP

OpenShift Container Platform requires a Google Cloud Platform (GCP) service account that provides authentication and authorization to access data in the Google APIs. If you do not have an existing IAM service account that contains the required roles in your project, you must create one.

Prerequisites

- You created a project to host your cluster.

Procedure


1. Create a service account in the project that you use to host your OpenShift Container Platform cluster. See [Creating a service account](#) in the GCP documentation.
2. Grant the service account the appropriate permissions. You can either grant the individual permissions that follow or assign the **Owner** role to it. See [Granting roles to a service account for specific resources](#).



NOTE

While making the service account an owner of the project is the easiest way to gain the required permissions, it means that service account has complete control over the project. You must determine if the risk that comes from offering that power is acceptable.

3. You can create the service account key in JSON format, or attach the service account to a GCP virtual machine. See [Creating service account keys](#) and [Creating and enabling service accounts for instances](#) in the GCP documentation.
You must have a service account key or a virtual machine with an attached service account to create the cluster.

**NOTE**

If you use a virtual machine with an attached service account to create your cluster, you must set **credentialsMode: Manual** in the **install-config.yaml** file before installation.

9.2.5.1. Required GCP roles

When you attach the **Owner** role to the service account that you create, you grant that service account all permissions, including those that are required to install OpenShift Container Platform. If your organization's security policies require a more restrictive set of permissions, you can create a service account with the following permissions. If you deploy your cluster into an existing virtual private cloud (VPC), the service account does not require certain networking permissions, which are noted in the following lists:

Required roles for the installation program

- Compute Admin
- Role Administrator
- Security Admin
- Service Account Admin
- Service Account Key Admin
- Service Account User
- Storage Admin

Required roles for creating network resources during installation

- DNS Administrator

Required roles for using the Cloud Credential Operator in passthrough mode

- Compute Load Balancer Admin

The following roles are applied to the service accounts that the control plane and compute machines use:

Table 9.4. GCP service account roles

Account	Roles
Control Plane	roles/compute.instanceAdmin
	roles/compute.networkAdmin
	roles/compute.securityAdmin
	roles/storage.admin

Account	Roles
	roles/iam.serviceAccountUser
Compute	roles/compute.viewer
	roles/storage.admin

9.2.5.2. Required GCP permissions for installer-provisioned infrastructure

When you attach the **Owner** role to the service account that you create, you grant that service account all permissions, including those that are required to install OpenShift Container Platform.

If your organization's security policies require a more restrictive set of permissions, you can create [custom roles](#) with the necessary permissions. The following permissions are required for the installer-provisioned infrastructure for creating and deleting the OpenShift Container Platform cluster.

Example 9.1. Required permissions for creating network resources

- **compute.addresses.create**
- **compute.addresses.createInternal**
- **compute.addresses.delete**
- **compute.addresses.get**
- **compute.addresses.list**
- **compute.addresses.use**
- **compute.addresses.useInternal**
- **compute.firewalls.create**
- **compute.firewalls.delete**
- **compute.firewalls.get**
- **compute.firewalls.list**
- **compute.forwardingRules.create**
- **compute.forwardingRules.get**
- **compute.forwardingRules.list**
- **compute.forwardingRules.setLabels**
- **compute.networks.create**
- **compute.networks.get**
- **compute.networks.list**

- **compute.networks.updatePolicy**
- **compute.routers.create**
- **compute.routers.get**
- **compute.routers.list**
- **compute.routers.update**
- **compute.routes.list**
- **compute.subnetworks.create**
- **compute.subnetworks.get**
- **compute.subnetworks.list**
- **compute.subnetworks.use**
- **compute.subnetworks.useExternallp**

Example 9.2. Required permissions for creating load balancer resources

- **compute.regionBackendServices.create**
- **compute.regionBackendServices.get**
- **compute.regionBackendServices.list**
- **compute.regionBackendServices.update**
- **compute.regionBackendServices.use**
- **compute.targetPools.addInstance**
- **compute.targetPools.create**
- **compute.targetPools.get**
- **compute.targetPools.list**
- **compute.targetPools.removeInstance**
- **compute.targetPools.use**

Example 9.3. Required permissions for creating DNS resources

- **dns.changes.create**
- **dns.changes.get**
- **dns.managedZones.create**
- **dns.managedZones.get**

- `dns.managedZones.list`
- `dns.networks.bindPrivateDNSZone`
- `dns.resourceRecordSets.create`
- `dns.resourceRecordSets.list`

Example 9.4. Required permissions for creating Service Account resources

- `iam.serviceAccountKeys.create`
- `iam.serviceAccountKeys.delete`
- `iam.serviceAccountKeys.get`
- `iam.serviceAccountKeys.list`
- `iam.serviceAccounts.actAs`
- `iam.serviceAccounts.create`
- `iam.serviceAccounts.delete`
- `iam.serviceAccounts.get`
- `iam.serviceAccounts.list`
- `resourcemanager.projects.get`
- `resourcemanager.projects.getIamPolicy`
- `resourcemanager.projects.setIamPolicy`

Example 9.5. Required permissions for creating compute resources

- `compute.disks.create`
- `compute.disks.get`
- `compute.disks.list`
- `compute.disks.setLabels`
- `compute.instanceGroups.create`
- `compute.instanceGroups.delete`
- `compute.instanceGroups.get`
- `compute.instanceGroups.list`
- `compute.instanceGroups.update`
- `compute.instanceGroups.use`

- **compute.instances.create**
- **compute.instances.delete**
- **compute.instances.get**
- **compute.instances.list**
- **compute.instances.setLabels**
- **compute.instances.setMetadata**
- **compute.instances.setServiceAccount**
- **compute.instances.setTags**
- **compute.instances.use**
- **compute.machineTypes.get**
- **compute.machineTypes.list**

Example 9.6. Required for creating storage resources

- **storage.buckets.create**
- **storage.buckets.delete**
- **storage.buckets.get**
- **storage.buckets.list**
- **storage.objects.create**
- **storage.objects.delete**
- **storage.objects.get**
- **storage.objects.list**

Example 9.7. Required permissions for creating health check resources

- **compute.healthChecks.create**
- **compute.healthChecks.get**
- **compute.healthChecks.list**
- **compute.healthChecks.useReadOnly**
- **compute.httpHealthChecks.create**
- **compute.httpHealthChecks.get**
- **compute.httpHealthChecks.list**

- `compute.httpHealthChecks.useReadOnly`

Example 9.8. Required permissions to get GCP zone and region related information

- `compute.globalOperations.get`
- `compute.regionOperations.get`
- `compute.regions.list`
- `compute.zoneOperations.get`
- `compute.zones.get`
- `compute.zones.list`

Example 9.9. Required permissions for checking services and quotas

- `monitoring.timeSeries.list`
- `serviceusage.quotas.get`
- `serviceusage.services.list`

Example 9.10. Required IAM permissions for installation

- `iam.roles.get`

Example 9.11. Optional Images permissions for installation

- `compute.images.list`

Example 9.12. Optional permission for running gather bootstrap

- `compute.instances.getSerialPortOutput`

Example 9.13. Required permissions for deleting network resources

- `compute.addresses.delete`
- `compute.addresses.deleteInternal`
- `compute.addresses.list`
- `compute.firewalls.delete`
- `compute.firewalls.list`
- `compute.forwardingRules.delete`

- **compute.forwardingRules.list**
- **compute.networks.delete**
- **compute.networks.list**
- **compute.networks.updatePolicy**
- **compute.routers.delete**
- **compute.routers.list**
- **compute.routes.list**
- **compute.subnetworks.delete**
- **compute.subnetworks.list**

Example 9.14. Required permissions for deleting load balancer resources

- **compute.regionBackendServices.delete**
- **compute.regionBackendServices.list**
- **compute.targetPools.delete**
- **compute.targetPools.list**

Example 9.15. Required permissions for deleting DNS resources

- **dns.changes.create**
- **dns.managedZones.delete**
- **dns.managedZones.get**
- **dns.managedZones.list**
- **dns.resourceRecordSets.delete**
- **dns.resourceRecordSets.list**

Example 9.16. Required permissions for deleting Service Account resources

- **iam.serviceAccounts.delete**
- **iam.serviceAccounts.get**
- **iam.serviceAccounts.list**
- **resourcemanager.projects.getIamPolicy**
- **resourcemanager.projects.setIamPolicy**

Example 9.17. Required permissions for deleting compute resources

- `compute.disks.delete`
- `compute.disks.list`
- `compute.instanceGroups.delete`
- `compute.instanceGroups.list`
- `compute.instances.delete`
- `compute.instances.list`
- `compute.instances.stop`
- `compute.machineTypes.list`

Example 9.18. Required for deleting storage resources

- `storage.buckets.delete`
- `storage.buckets.getIamPolicy`
- `storage.buckets.list`
- `storage.objects.delete`
- `storage.objects.list`

Example 9.19. Required permissions for deleting health check resources

- `compute.healthChecks.delete`
- `compute.healthChecks.list`
- `compute.httpHealthChecks.delete`
- `compute.httpHealthChecks.list`


Example 9.20. Required Images permissions for deletion

- `compute.images.list`

9.2.5.3. Required GCP permissions for shared VPC installations

When you are installing a cluster to a [shared VPC](#), you must configure the service account for both the host project and the service project. If you are not installing to a shared VPC, you can skip this section.

You must apply the minimum roles required for a standard installation as listed above, to the service project.



IMPORTANT

You can use granular permissions for a Cloud Credential Operator that operates in either manual or mint credentials mode. You cannot use granular permissions in passthrough credentials mode.

Ensure that the host project applies one of the following configurations to the service account:

Example 9.21. Required permissions for creating firewalls in the host project

- **projects/<host-project>/roles/dns.networks.bindPrivateDNSZone**
- **roles/compute.networkAdmin**
- **roles/compute.securityAdmin**

Example 9.22. Required minimal permissions

- **projects/<host-project>/roles/dns.networks.bindPrivateDNSZone**
- **roles/compute.networkUser**

9.2.6. Supported GCP regions

You can deploy an OpenShift Container Platform cluster to the following Google Cloud Platform (GCP) regions:

- **africa-south1** (Johannesburg, South Africa)
- **asia-east1** (Changhua County, Taiwan)
- **asia-east2** (Hong Kong)
- **asia-northeast1** (Tokyo, Japan)
- **asia-northeast2** (Osaka, Japan)
- **asia-northeast3** (Seoul, South Korea)
- **asia-south1** (Mumbai, India)
- **asia-south2** (Delhi, India)
- **asia-southeast1** (Jurong West, Singapore)
- **asia-southeast2** (Jakarta, Indonesia)
- **australia-southeast1** (Sydney, Australia)
- **australia-southeast2** (Melbourne, Australia)

- **europa-central2** (Warsaw, Poland)
- **europa-north1** (Hamina, Finland)
- **europa-southwest1** (Madrid, Spain)
- **europa-west1** (St. Ghislain, Belgium)
- **europa-west2** (London, England, UK)
- **europa-west3** (Frankfurt, Germany)
- **europa-west4** (Eemshaven, Netherlands)
- **europa-west6** (Zürich, Switzerland)
- **europa-west8** (Milan, Italy)
- **europa-west9** (Paris, France)
- **europa-west12** (Turin, Italy)
- **me-central1** (Doha, Qatar, Middle East)
- **me-central2** (Dammam, Saudi Arabia, Middle East)
- **me-west1** (Tel Aviv, Israel)
- **northamerica-northeast1** (Montréal, Québec, Canada)
- **northamerica-northeast2** (Toronto, Ontario, Canada)
- **southamerica-east1** (São Paulo, Brazil)
- **southamerica-west1** (Santiago, Chile)
- **us-central1** (Council Bluffs, Iowa, USA)
- **us-east1** (Moncks Corner, South Carolina, USA)
- **us-east4** (Ashburn, Northern Virginia, USA)
- **us-east5** (Columbus, Ohio)
- **us-south1** (Dallas, Texas)
- **us-west1** (The Dalles, Oregon, USA)
- **us-west2** (Los Angeles, California, USA)
- **us-west3** (Salt Lake City, Utah, USA)
- **us-west4** (Las Vegas, Nevada, USA)



NOTE

To determine which machine type instances are available by region and zone, see the Google [documentation](#).

9.2.7. Next steps

- Install an OpenShift Container Platform cluster on GCP. You can [install a customized cluster](#) or [quickly install a cluster](#) with default options.

9.3. INSTALLING A CLUSTER QUICKLY ON GCP

In OpenShift Container Platform version 4.16, you can install a cluster on Google Cloud Platform (GCP) that uses the default configuration options.

9.3.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [configured a GCP project](#) to host the cluster.
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.

9.3.2. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.



IMPORTANT

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

9.3.3. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the `core` user on each node, which enables password-less authentication.


After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The **./openshift-install gather** command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.



NOTE


You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.



NOTE

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the **./openshift-install gather** command.

**NOTE**


On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```

**NOTE**

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

9.3.4. Obtaining the installation program

Before you install OpenShift Container Platform, download the installation file on the host you are using for installation.


Prerequisites

- You have a computer that runs Linux or macOS, with 500 MB of local disk space.

Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select your infrastructure provider.

3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.



IMPORTANT

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.



IMPORTANT

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.

4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#). This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

9.3.5. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.

Procedure

1. Remove any existing GCP credentials that do not use the service account key for the GCP account that you configured for your cluster and that are stored in the following locations:

- The **GOOGLE_CREDENTIALS**, **GOOGLE_CLOUD_KEYFILE_JSON**, or **GKLOUD_KEYFILE_JSON** environment variables
 - The `~/gcp/osServiceAccount.json` file
 - The **gcloud cli** default credentials
2. Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

- 1** For `<installation_directory>`, specify the directory name to store the files that the installation program creates.
- 2** To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

When specifying the directory:

- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
 - Use an empty directory. Some installation assets, such as bootstrap X.509 certificates, have short expiration intervals, therefore you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.
3. Provide values at the prompts:
- Optional: Select an SSH key to use to access your cluster machines.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- Select **gcp** as the platform to target.
- If you have not configured the service account key for your GCP account on your host, you must obtain it from GCP and paste the contents of the file or enter the absolute path to the file.
- Select the project ID to provision the cluster in. The default value is specified by the service account that you configured.
- Select the region to deploy the cluster to.
- Select the base domain to deploy the cluster to. The base domain corresponds to the public DNS zone that you created for your cluster.

- g. Enter a descriptive name for your cluster. If you provide a name that is longer than 6 characters, only the first 6 characters will be used in the infrastructure ID that is generated from the cluster name.
 - h. Paste the [pull secret from Red Hat OpenShift Cluster Manager](#) .
4. Optional: You can reduce the number of permissions for the service account that you used to install the cluster.
- If you assigned the **Owner** role to your service account, you can remove that role and replace it with the **Viewer** role.
 - If you included the **Service Account Key Admin** role, you can remove it.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to `<installation_directory>/openshift_install.log`.



IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

9.3.6. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.



IMPORTANT

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**.
To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.

To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:


```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.



NOTE

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your **PATH**.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

9.3.7. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

Additional resources

- See [Accessing the web console](#) for more details about accessing and understanding the OpenShift Container Platform web console.

9.3.8. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service

9.3.9. Next steps

- [Customize your cluster](#).
- If necessary, you can [opt out of remote health reporting](#).

9.4. INSTALLING A CLUSTER ON GCP WITH CUSTOMIZATIONS

In OpenShift Container Platform version 4.16, you can install a customized cluster on infrastructure that the installation program provisions on Google Cloud Platform (GCP). To customize the installation, you modify parameters in the **install-config.yaml** file before you install the cluster.

9.4.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.


- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [configured a GCP project](#) to host the cluster.
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.

9.4.2. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.



IMPORTANT


If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

9.4.3. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.



NOTE

You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.



NOTE

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.



NOTE

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```



NOTE

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

9.4.4. Obtaining the installation program

Before you install OpenShift Container Platform, download the installation file on the host you are using for installation.

Prerequisites

- You have a computer that runs Linux or macOS, with 500 MB of local disk space.

Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select your infrastructure provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.



IMPORTANT

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.



IMPORTANT

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.

4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#) . This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

9.4.5. Creating the installation configuration file

You can customize the OpenShift Container Platform cluster you install on Google Cloud Platform (GCP).

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create the **install-config.yaml** file.
 - a. Change to the directory that contains the installation program and run the following command:

```
└─$ ./openshift-install create install-config --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the directory name to store the files that the installation program creates.

When specifying the directory:

- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
 - Use an empty directory. Some installation assets, such as bootstrap X.509 certificates, have short expiration intervals, therefore you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.
- b. At the prompts, provide the configuration details for your cloud:
 - i. Optional: Select an SSH key to use to access your cluster machines.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- ii. Select **gcp** as the platform to target.
- iii. If you have not configured the service account key for your GCP account on your computer, you must obtain it from GCP and paste the contents of the file or enter the absolute path to the file.

- iv. Select the project ID to provision the cluster in. The default value is specified by the service account that you configured.
 - v. Select the region to deploy the cluster to.
 - vi. Select the base domain to deploy the cluster to. The base domain corresponds to the public DNS zone that you created for your cluster.
 - vii. Enter a descriptive name for your cluster.
2. Modify the **install-config.yaml** file. You can find more information about the available parameters in the "Installation configuration parameters" section.

**NOTE**

If you are installing a three-node cluster, be sure to set the **compute.replicas** parameter to **0**. This ensures that the cluster's control planes are schedulable. For more information, see "Installing a three-node cluster on GCP".

3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.

**IMPORTANT**

The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

Additional resources

- [Installation configuration parameters for GCP](#)

9.4.5.1. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 9.5. Minimum resource requirements

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Bootstrap	RHCOS	4	16 GB	100 GB	300
Control plane	RHCOS	4	16 GB	100 GB	300
Compute	RHCOS, RHEL 8.6 and later [3]	2	8 GB	100 GB	300

1. One vCPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio: (threads per core × cores) × sockets = vCPUs.

2. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes which require a 10 ms p99 fsync duration. Note that on many cloud platforms, storage size and IOPS scale together, so you might need to over-allocate storage volume to obtain sufficient performance.
3. As with all user-provisioned installations, if you choose to use RHEL compute machines in your cluster, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Use of RHEL 7 compute machines is deprecated and has been removed in OpenShift Container Platform 4.10 and later.



NOTE

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

9.4.5.2. Tested instance types for GCP

The following Google Cloud Platform instance types have been tested with OpenShift Container Platform.

Example 9.23. Machine series

- **A2**
- **A3**
- **C2**
- **C2D**
- **C3**
- **C3D**
- **E2**
- **M1**

- N1
- N2
- N2D
- Tau T2D

9.4.5.3. Tested instance types for GCP on 64-bit ARM infrastructures

The following Google Cloud Platform (GCP) 64-bit ARM instance types have been tested with OpenShift Container Platform.

Example 9.24. Machine series for 64-bit ARM machines

- Tau T2A

9.4.5.4. Using custom machine types

Using a custom machine type to install a OpenShift Container Platform cluster is supported.

Consider the following when using a custom machine type:

- Similar to predefined instance types, custom machine types must meet the minimum resource requirements for control plane and compute machines. For more information, see "Minimum resource requirements for cluster installation".
- The name of the custom machine type must adhere to the following syntax:
custom-<number_of_cpus>-<amount_of_memory_in_mb>

For example, **custom-6-20480**.

As part of the installation process, you specify the custom machine type in the **install-config.yaml** file.

Sample install-config.yaml file with a custom machine type

```
compute:
- architecture: amd64
  hyperthreading: Enabled
  name: worker
  platform:
    gcp:
      type: custom-6-20480
  replicas: 2
controlPlane:
  architecture: amd64
  hyperthreading: Enabled
  name: master
  platform:
    gcp:
      type: custom-6-20480
  replicas: 3
```

9.4.5.5. Enabling Shielded VMs

You can use Shielded VMs when installing your cluster. Shielded VMs have extra security features including secure boot, firmware and integrity monitoring, and rootkit detection. For more information, see Google's documentation on [Shielded VMs](#).



NOTE

Shielded VMs are currently not supported on clusters with 64-bit ARM infrastructures.

Prerequisites

- You have created an **install-config.yaml** file.

Procedure

- Use a text editor to edit the **install-config.yaml** file prior to deploying your cluster and add one of the following stanzas:
 - To use shielded VMs for only control plane machines:

```
controlPlane:
  platform:
    gcp:
      secureBoot: Enabled
```

- To use shielded VMs for only compute machines:

```
compute:
- platform:
  gcp:
    secureBoot: Enabled
```

- To use shielded VMs for all machines:

```
platform:
  gcp:
    defaultMachinePlatform:
      secureBoot: Enabled
```

9.4.5.6. Enabling Confidential VMs

You can use Confidential VMs when installing your cluster. Confidential VMs encrypt data while it is being processed. For more information, see Google's documentation on [Confidential Computing](#). You can enable Confidential VMs and Shielded VMs at the same time, although they are not dependent on each other.



NOTE

Confidential VMs are currently not supported on 64-bit ARM architectures.

Prerequisites

- You have created an **install-config.yaml** file.

Procedure

- Use a text editor to edit the **install-config.yaml** file prior to deploying your cluster and add one of the following stanzas:
 - To use confidential VMs for only control plane machines:

```
controlPlane:
  platform:
    gcp:
      confidentialCompute: Enabled 1
      type: n2d-standard-8 2
      onHostMaintenance: Terminate 3
```

- Enable confidential VMs.
- Specify a machine type that supports Confidential VMs. Confidential VMs require the N2D or C2D series of machine types. For more information on supported machine types, see [Supported operating systems and machine types](#).
- Specify the behavior of the VM during a host maintenance event, such as a hardware or software update. For a machine that uses Confidential VM, this value must be set to **Terminate**, which stops the VM. Confidential VMs do not support live VM migration.

- To use confidential VMs for only compute machines:

```
compute:
  - platform:
      gcp:
        confidentialCompute: Enabled
        type: n2d-standard-8
        onHostMaintenance: Terminate
```

- To use confidential VMs for all machines:

```
platform:
  gcp:
    defaultMachinePlatform:
      confidentialCompute: Enabled
      type: n2d-standard-8
      onHostMaintenance: Terminate
```

9.4.5.7. Sample customized install-config.yaml file for GCP

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.



IMPORTANT

This sample YAML file is provided for reference only. You must obtain your **install-config.yaml** file by using the installation program and modify it.

```
apiVersion: v1
baseDomain: example.com 1
credentialsMode: Mint 2
controlPlane: 3 4
  hyperthreading: Enabled 5
  name: master
  platform:
    gcp:
      type: n2-standard-4
      zones:
      - us-central1-a
      - us-central1-c
    osDisk:
      diskType: pd-ssd
      diskSizeGB: 1024
      encryptionKey: 6
      kmsKey:
        name: worker-key
        keyRing: test-machine-keys
        location: global
        projectID: project-id
    tags: 7
    - control-plane-tag1
    - control-plane-tag2
    osImage: 8
      project: example-project-name
      name: example-image-name
  replicas: 3
compute: 9 10
- hyperthreading: Enabled 11
  name: worker
  platform:
    gcp:
      type: n2-standard-4
      zones:
      - us-central1-a
      - us-central1-c
    osDisk:
      diskType: pd-standard
      diskSizeGB: 128
      encryptionKey: 12
      kmsKey:
        name: worker-key
        keyRing: test-machine-keys
        location: global
        projectID: project-id
    tags: 13
    - compute-tag1
    - compute-tag2
    osImage: 14
      project: example-project-name
      name: example-image-name
  replicas: 3
metadata:
```

```

name: test-cluster 15
networking:
  clusterNetwork:
    - cidr: 10.128.0.0/14
      hostPrefix: 23
  machineNetwork:
    - cidr: 10.0.0.0/16
  networkType: OVNKubernetes 16
  serviceNetwork:
    - 172.30.0.0/16
platform:
  gcp:
    projectID: openshift-production 17
    region: us-central1 18
    defaultMachinePlatform:
      tags: 19
      - global-tag1
      - global-tag2
    osImage: 20
      project: example-project-name
      name: example-image-name
  pullSecret: '{"auths": ...}' 21
  fips: false 22
  sshKey: ssh-ed25519 AAAA... 23

```

- 1 15 17 18 21** Required. The installation program prompts you for this value.
- 2** Optional: Add this parameter to force the Cloud Credential Operator (CCO) to use the specified mode. By default, the CCO uses the root credentials in the **kube-system** namespace to dynamically try to determine the capabilities of the credentials. For details about CCO modes, see the "About the Cloud Credential Operator" section in the *Authentication and authorization* guide.
- 3 9** If you do not provide these parameters and values, the installation program provides the default value.
- 4 10** The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.
- 5 11** Whether to enable or disable simultaneous multithreading, or **hyperthreading**. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores. You can disable it by setting the parameter value to **Disabled**. If you disable simultaneous multithreading in some cluster machines, you must disable it in all cluster machines.



IMPORTANT

If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance. Use larger machine types, such as **n1-standard-8**, for your machines if you disable simultaneous multithreading.

- 6 12** Optional: The custom encryption key section to encrypt both virtual machines and persistent

- 7 13 19 Optional: A set of network tags to apply to the control plane or compute machine sets. The **platform.gcp.defaultMachinePlatform.tags** parameter will apply to both control plane and
- 8 14 20 Optional: A custom Red Hat Enterprise Linux CoreOS (RHCOS) that should be used to boot control plane and compute machines. The **project** and **name** parameters under **platform.gcp.defaultMachinePlatform.osImage** apply to both control plane and compute machines. If the **project** and **name** parameters under **controlPlane.platform.gcp.osImage** or **compute.platform.gcp.osImage** are set, they override the **platform.gcp.defaultMachinePlatform.osImage** parameters.
- 16 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 22 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.



IMPORTANT

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 23 You can optionally provide the **sshKey** value that you use to access the machines in your cluster.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

Additional resources

- [Enabling customer-managed encryption keys for a compute machine set](#)

9.4.5.8. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.



NOTE

The installation program does not support the proxy **readinessEndpoints** field.

**NOTE**

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.

**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

9.4.6. Managing user-defined labels and tags for GCP

**IMPORTANT**

Support for user-defined labels and tags for GCP is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

Google Cloud Platform (GCP) provides labels and tags that help to identify and organize the resources created for a specific OpenShift Container Platform cluster, making them easier to manage.

You can define labels and tags for each GCP resource only during OpenShift Container Platform cluster installation.

**IMPORTANT**

User-defined labels and tags are not supported for OpenShift Container Platform clusters upgraded to OpenShift Container Platform 4.16.

**NOTE**

You cannot update the tags that are already added. Also, a new tag-supported resource creation fails if the configured tag keys or tag values are deleted.

User-defined labels

User-defined labels and OpenShift Container Platform specific labels are applied only to resources created by OpenShift Container Platform installation program and its core components such as:

- GCP filestore CSI Driver Operator
- GCP PD CSI Driver Operator
- Image Registry Operator
- Machine API provider for GCP

User-defined tags are attached to resources created by the OpenShift Container Platform installation program, Image Registry Operator, and Machine API Operator. User-defined tags are not attached to the resources created by any other Operators or the Kubernetes in-tree components.

User-defined labels and OpenShift Container Platform labels are available on the following GCP resources:

- Compute disk
- Compute instance
- Compute image
- Compute forwarding rule
- DNS managed zone
- Filestore instance
- Storage bucket

Limitations to user-defined labels

- Labels for **ComputeAddress** are supported in the GCP beta version. OpenShift Container Platform does not add labels to the resource.

User-defined tags

User-defined tags are attached to resources created by the OpenShift Container Platform Image Registry Operator and not on the resources created by any other Operators or the Kubernetes in-tree components.

User-defined tags are available on the following GCP resources:

- Storage buckets
- Compute instances
- Compute disks

Limitations to the user-defined tags

- Tags will not be attached to the following items:
 - Filestore instance resources created by the GCP filestore CSI driver Operator
 - Compute disk and compute image resources created by the GCP PD CSI driver Operator
- Tags must not be restricted to particular service accounts, because Operators create and use service accounts with minimal roles.

- OpenShift Container Platform does not create any key and value resources of the tag.
- OpenShift Container Platform specific tags are not added to any resource.

Additional resources

- For more information about identifying the **OrganizationID**, see: [OrganizationID](#)
- For more information about identifying the **ProjectID**, see: [ProjectID](#)
- For more information about labels, see [Labels Overview](#).
- For more information about tags, see [Tags Overview](#).


9.4.6.1. Configuring user-defined labels and tags for GCP

Prerequisites

- The installation program requires that a service account includes a **TagUser** role, so that the program can create the OpenShift Container Platform cluster with defined tags at both organization and project levels.

Procedure

- Update the **install-config.yaml** file to define the list of desired labels and tags.



NOTE

Labels and tags are defined during the **install-config.yaml** creation phase, and cannot be modified or updated with new labels and tags after cluster creation.

Sample install-config.yaml file

```
apiVersion: v1
featureSet: TechPreviewNoUpgrade
platform:
  gcp:
    userLabels: 1
    - key: <label_key> 2
      value: <label_value> 3
    userTags: 4
    - parentID: <OrganizationID/ProjectID> 5
      key: <tag_key_short_name>
      value: <tag_value_short_name>
```

- 1 Adds keys and values as labels to the resources created on GCP.
- 2 Defines the label name.
- 3 Defines the label content.
- 4 Adds keys and values as tags to the resources created on GCP.

- 5 The ID of the hierarchical resource where the tags are defined, at the organization or the project level.

The following are the requirements for user-defined labels:

- A label key and value must have a minimum of 1 character and can have a maximum of 63 characters.
- A label key and value must contain only lowercase letters, numeric characters, underscore (`_`), and dash (`-`).
- A label key must start with a lowercase letter.
- You can configure a maximum of 32 labels per resource. Each resource can have a maximum of 64 labels, and 32 labels are reserved for internal use by OpenShift Container Platform.

The following are the requirements for user-defined tags:

- Tag key and tag value must already exist. OpenShift Container Platform does not create the key and the value.
- A tag **parentID** can be either **OrganizationID** or **ProjectID**:
 - **OrganizationID** must consist of decimal numbers without leading zeros.
 - **ProjectID** must be 6 to 30 characters in length, that includes only lowercase letters, numbers, and hyphens.
 - **ProjectID** must start with a letter, and cannot end with a hyphen.
- A tag key must contain only uppercase and lowercase alphanumeric characters, hyphen (`-`), underscore (`_`), and period (`.`).
- A tag value must contain only uppercase and lowercase alphanumeric characters, hyphen (`-`), underscore (`_`), period (`.`), at sign (`@`), percent sign (`%`), equals sign (`=`), plus (`+`), colon (`:`), comma (`,`), asterisk (`*`), pound sign (`$`), ampersand (`&`), parentheses (`()`), square braces (`[]`), curly braces (`{}`), and space.
- A tag key and value must begin and end with an alphanumeric character.
- Tag value must be one of the pre-defined values for the key.
- You can configure a maximum of 50 tags.
- There should be no tag key defined with the same value as any of the existing tag keys that will be inherited from the parent resource.

9.4.6.2. Querying user-defined labels and tags for GCP

After creating the OpenShift Container Platform cluster, you can access the list of the labels and tags defined for the GCP resources in the **infrastructures.config.openshift.io/cluster** object as shown in the following sample **infrastructure.yaml** file.

Sample **infrastructure.yaml** file

```
apiVersion: config.openshift.io/v1
```

```

kind: Infrastructure
metadata:
  name: cluster
spec:
  platformSpec:
    type: GCP
status:
  infrastructureName: <cluster_id> 1
  platform: GCP
  platformStatus:
    gcp:
      resourceLabels:
        - key: <label_key>
          value: <label_value>
      resourceTags:
        - key: <tag_key_short_name>
          parentID: <OrganizationID/ProjectID>
          value: <tag_value_short_name>
    type: GCP

```

- 1** The cluster ID that is generated during cluster installation.

Along with the user-defined labels, resources have a label defined by the OpenShift Container Platform. The format of the OpenShift Container Platform labels is **kubernetes-io-cluster-`<cluster_id>`:owned**.

9.4.7. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.



IMPORTANT

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**.

To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.



NOTE

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.

5. Move the **oc** binary to a directory on your PATH.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

9.4.8. Alternatives to storing administrator-level secrets in the kube-system project

By default, administrator secrets are stored in the **kube-system** project. If you configured the **credentialsMode** parameter in the **install-config.yaml** file to **Manual**, you must use one of the following alternatives:

- To manage long-term cloud credentials manually, follow the procedure in [Manually creating long-term credentials](#).
- To implement short-term credentials that are managed outside the cluster for individual components, follow the procedures in [Configuring a GCP cluster to use short-term credentials](#).

9.4.8.1. Manually creating long-term credentials

The Cloud Credential Operator (CCO) can be put into manual mode prior to installation in environments where the cloud identity and access management (IAM) APIs are not reachable, or the administrator prefers not to store an administrator-level credential secret in the cluster **kube-system** namespace.

Procedure

1. Add the following granular permissions to the GCP account that the installation program uses:

Example 9.25. Required GCP permissions

- compute.machineTypes.list
- compute.regions.list
- compute.zones.list
- dns.changes.create
- dns.changes.get
- dns.managedZones.create
- dns.managedZones.delete
- dns.managedZones.get
- dns.managedZones.list
- dns.networks.bindPrivateDNSZone

- `dns.resourceRecordSets.create`
- `dns.resourceRecordSets.delete`
- `dns.resourceRecordSets.list`

- If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```
apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...
```

- If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

- Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

- Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included \ 1
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \ 2
  --to=<path_to_directory_for_credentials_requests> \ 3
```

- 1 The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2 Specify the location of the **install-config.yaml** file.
- 3 Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

This command creates a YAML file for each **CredentialsRequest** object.

Sample CredentialsRequest object

```
apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
```

```

metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: GCPProviderSpec
    predefinedRoles:
      - roles/storage.admin
      - roles/iam.serviceAccountUser
    skipServiceCheck: true
  ...

```

6. Create YAML files for secrets in the **openshift-install** manifests directory that you generated previously. The secrets must be stored using the namespace and secret name defined in the **spec.secretRef** for each **CredentialsRequest** object.

Sample CredentialsRequest object with secrets

```

apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    ...
  secretRef:
    name: <component_secret>
    namespace: <component_namespace>
  ...

```

Sample Secret object

```

apiVersion: v1
kind: Secret
metadata:
  name: <component_secret>
  namespace: <component_namespace>
data:
  service_account.json: <base64_encoded_gcp_service_account_file>

```



IMPORTANT


Before upgrading a cluster that uses manually maintained credentials, you must ensure that the CCO is in an upgradeable state.

9.4.8.2. Configuring a GCP cluster to use short-term credentials

To install a cluster that is configured to use GCP Workload Identity, you must configure the CCO utility and create the required GCP resources for your cluster.

9.4.8.2.1. Configuring the Cloud Credential Operator utility

To create and manage cloud credentials from outside of the cluster when the Cloud Credential Operator (CCO) is operating in manual mode, extract and prepare the CCO utility (**ccoctl**) binary.



NOTE

The **ccoctl** utility is a Linux binary that must run in a Linux environment.

Prerequisites

- You have access to an OpenShift Container Platform account with cluster administrator access.
- You have installed the OpenShift CLI (**oc**).
- You have added one of the following authentication options to the GCP account that the installation program uses:
 - The **IAM Workload Identity Pool Admin** role.
 - The following granular permissions:

Example 9.26. Required GCP permissions

- `compute.projects.get`
- `iam.googleapis.com/workloadIdentityPoolProviders.create`
- `iam.googleapis.com/workloadIdentityPoolProviders.get`
- `iam.googleapis.com/workloadIdentityPools.create`
- `iam.googleapis.com/workloadIdentityPools.delete`
- `iam.googleapis.com/workloadIdentityPools.get`
- `iam.googleapis.com/workloadIdentityPools.undelete`
- `iam.roles.create`
- `iam.roles.delete`
- `iam.roles.list`
- `iam.roles.undelete`
- `iam.roles.update`
- `iam.serviceAccounts.create`
- `iam.serviceAccounts.delete`
- `iam.serviceAccounts.getIamPolicy`
- `iam.serviceAccounts.list`
- `iam.serviceAccounts.setIamPolicy`

- iam.workloadIdentityPoolProviders.get
- iam.workloadIdentityPools.delete
- resourcemanager.projects.get
- resourcemanager.projects.getIamPolicy
- resourcemanager.projects.setIamPolicy
- storage.buckets.create
- storage.buckets.delete
- storage.buckets.get
- storage.buckets.getIamPolicy
- storage.buckets.setIamPolicy
- storage.objects.create
- storage.objects.delete
- storage.objects.list


Procedure

1. Set a variable for the OpenShift Container Platform release image by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

2. Obtain the CCO container image from the OpenShift Container Platform release image by running the following command:

```
$ CCO_IMAGE=$(oc adm release info --image-for='cloud-credential-operator'
$RELEASE_IMAGE -a ~/.pull-secret)
```



NOTE

Ensure that the architecture of the **\$RELEASE_IMAGE** matches the architecture of the environment in which you will use the **ccoctl** tool.

3. Extract the **ccoctl** binary from the CCO container image within the OpenShift Container Platform release image by running the following command:

```
$ oc image extract $CCO_IMAGE \
--file="/usr/bin/ccoctl.<rhel_version>" \
-a ~/.pull-secret
```

- 1 For **<rhel_version>**, specify the value that corresponds to the version of Red Hat Enterprise Linux (RHEL) that the host uses. If no value is specified, **ccoctl.rhel8** is used by default. The following values are valid:

- **rhel8**: Specify this value for hosts that use RHEL 8.
- **rhel9**: Specify this value for hosts that use RHEL 9.

4. Change the permissions to make **ccoctl** executable by running the following command:

```
$ chmod 775 ccoctl.<rhel_version>
```

Verification

- To verify that **ccoctl** is ready to use, display the help file by running the following command:

```
$ ccoctl --help
```

Output of **ccoctl --help**

```
OpenShift credentials provisioning tool
```

```
Usage:
```

```
ccoctl [command]
```

```
Available Commands:
```

```
aws      Manage credentials objects for AWS cloud
azure    Manage credentials objects for Azure
gcp      Manage credentials objects for Google cloud
help     Help about any command
ibmcloud Manage credentials objects for IBM Cloud
nutanix  Manage credentials objects for Nutanix
```

```
Flags:
```

```
-h, --help  help for ccoctl
```

```
Use "ccoctl [command] --help" for more information about a command.
```

9.4.8.2.2. Creating GCP resources with the Cloud Credential Operator utility

You can use the **ccoctl gcp create-all** command to automate the creation of GCP resources.



NOTE

By default, **ccoctl** creates objects in the directory in which the commands are run. To create the objects in a different directory, use the **--output-dir** flag. This procedure uses **<path_to_ccoctl_output_dir>** to refer to this directory.

Prerequisites

You must have:

- Extracted and prepared the **ccoctl** binary.

Procedure

1. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

2. Extract the list of **CredentialsRequest** objects from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included \1
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \2
  --to=<path_to_directory_for_credentials_requests> \3
```

- 1 The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2 Specify the location of the **install-config.yaml** file.
- 3 Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.



NOTE

This command might take a few moments to run.

3. Use the **ccoctl** tool to process all **CredentialsRequest** objects by running the following command:

```
$ ccoctl gcp create-all \
  --name=<name> \1
  --region=<gcp_region> \2
  --project=<gcp_project_id> \3
  --credentials-requests-dir=<path_to_credentials_requests_directory> \4
```

- 1 Specify the user-defined name for all created GCP resources used for tracking.
- 2 Specify the GCP region in which cloud resources will be created.
- 3 Specify the GCP project ID in which cloud resources will be created.
- 4 Specify the directory containing the files of **CredentialsRequest** manifests to create GCP service accounts.



NOTE

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

Verification

- To verify that the OpenShift Container Platform secrets are created, list the files in the `<path_to_ccoctl_output_dir>/manifests` directory:

```
$ ls <path_to_ccoctl_output_dir>/manifests
```

Example output

```
cluster-authentication-02-config.yaml
openshift-cloud-controller-manager-gcp-ccm-cloud-credentials-credentials.yaml
openshift-cloud-credential-operator-cloud-credential-operator-gcp-ro-creds-credentials.yaml
openshift-cloud-network-config-controller-cloud-credentials-credentials.yaml
openshift-cluster-api-capg-manager-bootstrap-credentials-credentials.yaml
openshift-cluster-csi-drivers-gcp-pd-cloud-credentials-credentials.yaml
openshift-image-registry-installer-cloud-credentials-credentials.yaml
openshift-ingress-operator-cloud-credentials-credentials.yaml
openshift-machine-api-gcp-cloud-credentials-credentials.yaml
```

You can verify that the IAM service accounts are created by querying GCP. For more information, refer to [GCP documentation on listing IAM service accounts](#).

9.4.8.2.3. Incorporating the Cloud Credential Operator utility manifests

To implement short-term security credentials managed outside the cluster for individual components, you must move the manifest files that the Cloud Credential Operator utility (**ccoctl**) created to the correct directories for the installation program.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have configured the Cloud Credential Operator utility (**ccoctl**).
- You have created the cloud provider resources that are required for your cluster with the **ccoctl** utility.

Procedure

1. Add the following granular permissions to the GCP account that the installation program uses:

Example 9.27. Required GCP permissions

- `compute.machineTypes.list`
- `compute.regions.list`
- `compute.zones.list`
- `dns.changes.create`
- `dns.changes.get`
- `dns.managedZones.create`
- `dns.managedZones.delete`

- `dns.managedZones.get`
- `dns.managedZones.list`
- `dns.networks.bindPrivateDNSZone`
- `dns.resourceRecordSets.create`
- `dns.resourceRecordSets.delete`
- `dns.resourceRecordSets.list`

2. If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```
apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...
```

3. If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

4. Copy the manifests that the **ccoctl** utility generated to the **manifests** directory that the installation program created by running the following command:

```
$ cp /<path_to_ccoctl_output_dir>/manifests/* ./manifests/
```

5. Copy the private key that the **ccoctl** utility generated in the **tls** directory to the installation directory by running the following command:

```
$ cp -a /<path_to_ccoctl_output_dir>/tls .
```

9.4.9. Using the GCP Marketplace offering

Using the GCP Marketplace offering lets you deploy an OpenShift Container Platform cluster, which is billed on pay-per-use basis (hourly, per core) through GCP, while still being supported directly by Red Hat.

By default, the installation program downloads and installs the Red Hat Enterprise Linux CoreOS (RHCOS) image that is used to deploy compute machines. To deploy an OpenShift Container Platform cluster using an RHCOS image from the GCP Marketplace, override the default behavior by modifying the **install-config.yaml** file to reference the location of GCP Marketplace offer.

Prerequisites

- You have an existing **install-config.yaml** file.

Procedure


1. Edit the **compute.platform.gcp.osImage** parameters to specify the location of the GCP Marketplace image:
 - Set the **project** parameter to **redhat-marketplace-public**
 - Set the **name** parameter to one of the following offers:
 - OpenShift Container Platform**
redhat-coreos-ocp-413-x86-64-202305021736
 - OpenShift Platform Plus**
redhat-coreos-opp-413-x86-64-202305021736
 - OpenShift Kubernetes Engine**
redhat-coreos-oke-413-x86-64-202305021736
2. Save the file and reference it when deploying the cluster.

Sample install-config.yaml file that specifies a GCP Marketplace image for compute machines

```
apiVersion: v1
baseDomain: example.com
controlPlane:
# ...
compute:
  platform:
    gcp:
      osImage:
        project: redhat-marketplace-public
        name: redhat-coreos-ocp-413-x86-64-202305021736
# ...
```

9.4.10. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.

Procedure

- Remove any existing GCP credentials that do not use the service account key for the GCP account that you configured for your cluster and that are stored in the following locations:
 - The **GOOGLE_CREDENTIALS**, **GOOGLE_CLOUD_KEYFILE_JSON**, or **GKLOUD_KEYFILE_JSON** environment variables
 - The `~/gcp/osServiceAccount.json` file
 - The **gcloud cli** default credentials
- Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

- For `<installation_directory>`, specify the location of your customized `./install-config.yaml` file.
- To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

- Optional: You can reduce the number of permissions for the service account that you used to install the cluster.
 - If you assigned the **Owner** role to your service account, you can remove that role and replace it with the **Viewer** role.
 - If you included the **Service Account Key Admin** role, you can remove it.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to `<installation_directory>/openshift_install.log`.



IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```




IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

9.4.11. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

Additional resources

- See [Accessing the web console](#) for more details about accessing and understanding the OpenShift Container Platform web console.

9.4.12. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service

9.4.13. Next steps

- [Customize your cluster](#).
- If necessary, you can [opt out of remote health reporting](#) .

9.5. INSTALLING A CLUSTER ON GCP WITH NETWORK CUSTOMIZATIONS

In OpenShift Container Platform version 4.16, you can install a cluster with a customized network configuration on infrastructure that the installation program provisions on Google Cloud Platform (GCP). By customizing your network configuration, your cluster can coexist with existing IP address allocations in your environment and integrate with existing MTU and VXLAN configurations. To customize the installation, you modify parameters in the **install-config.yaml** file before you install the cluster.

You must set most of the network configuration parameters during installation, and you can modify only **kubeProxy** configuration parameters in a running cluster.

9.5.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [configured a GCP project](#) to host the cluster.
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.

9.5.2. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.

- Obtain the packages that are required to perform cluster updates.



IMPORTANT

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

9.5.3. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.


After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.



NOTE

You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.

**NOTE**

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.

**NOTE**

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```

**NOTE**

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

9.5.4. Obtaining the installation program


Before you install OpenShift Container Platform, download the installation file on the host you are using for installation.

Prerequisites

- You have a computer that runs Linux or macOS, with 500 MB of local disk space.


Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select your infrastructure provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.



IMPORTANT

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.



IMPORTANT

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.

4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#) . This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

9.5.5. Creating the installation configuration file

You can customize the OpenShift Container Platform cluster you install on Google Cloud Platform (GCP).

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create the **install-config.yaml** file.
 - a. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install create install-config --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the directory name to store the files that the installation program creates.

When specifying the directory:

- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
 - Use an empty directory. Some installation assets, such as bootstrap X.509 certificates, have short expiration intervals, therefore you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.
- b. At the prompts, provide the configuration details for your cloud:
 - i. Optional: Select an SSH key to use to access your cluster machines.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- ii. Select **gcp** as the platform to target.
 - iii. If you have not configured the service account key for your GCP account on your computer, you must obtain it from GCP and paste the contents of the file or enter the absolute path to the file.
 - iv. Select the project ID to provision the cluster in. The default value is specified by the service account that you configured.
 - v. Select the region to deploy the cluster to.
 - vi. Select the base domain to deploy the cluster to. The base domain corresponds to the public DNS zone that you created for your cluster.
 - vii. Enter a descriptive name for your cluster.
2. Modify the **install-config.yaml** file. You can find more information about the available parameters in the "Installation configuration parameters" section.

3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

Additional resources

- [Installation configuration parameters for GCP](#)

9.5.5.1. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 9.6. Minimum resource requirements

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Bootstrap	RHCOS	4	16 GB	100 GB	300
Control plane	RHCOS	4	16 GB	100 GB	300
Compute	RHCOS, RHEL 8.6 and later [3]	2	8 GB	100 GB	300

1. One vCPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio: $(\text{threads per core} \times \text{cores}) \times \text{sockets} = \text{vCPUs}$.
2. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes which require a 10 ms p99 fsync duration. Note that on many cloud platforms, storage size and IOPS scale together, so you might need to over-allocate storage volume to obtain sufficient performance.
3. As with all user-provisioned installations, if you choose to use RHEL compute machines in your cluster, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Use of RHEL 7 compute machines is deprecated and has been removed in OpenShift Container Platform 4.10 and later.

**NOTE**

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

9.5.5.2. Tested instance types for GCP

The following Google Cloud Platform instance types have been tested with OpenShift Container Platform.

Example 9.28. Machine series

- **A2**
- **A3**
- **C2**
- **C2D**
- **C3**
- **C3D**
- **E2**
- **M1**
- **N1**
- **N2**
- **N2D**
- **Tau T2D**

9.5.5.3. Tested instance types for GCP on 64-bit ARM infrastructures

The following Google Cloud Platform (GCP) 64-bit ARM instance types have been tested with OpenShift Container Platform.

Example 9.29. Machine series for 64-bit ARM machines

- **Tau T2A**

9.5.5.4. Using custom machine types

Using a custom machine type to install a OpenShift Container Platform cluster is supported.

Consider the following when using a custom machine type:

- Similar to predefined instance types, custom machine types must meet the minimum resource requirements for control plane and compute machines. For more information, see "Minimum resource requirements for cluster installation".
- The name of the custom machine type must adhere to the following syntax:
custom-<number_of_cpus>-<amount_of_memory_in_mb>

For example, **custom-6-20480**.

As part of the installation process, you specify the custom machine type in the **install-config.yaml** file.

Sample install-config.yaml file with a custom machine type

```
compute:
- architecture: amd64
  hyperthreading: Enabled
  name: worker
  platform:
    gcp:
      type: custom-6-20480
  replicas: 2
controlPlane:
  architecture: amd64
  hyperthreading: Enabled
  name: master
  platform:
    gcp:
      type: custom-6-20480
  replicas: 3
```

9.5.5.5. Enabling Shielded VMs

You can use Shielded VMs when installing your cluster. Shielded VMs have extra security features including secure boot, firmware and integrity monitoring, and rootkit detection. For more information, see Google's documentation on [Shielded VMs](#).



NOTE

Shielded VMs are currently not supported on clusters with 64-bit ARM infrastructures.

Prerequisites

- You have created an **install-config.yaml** file.

Procedure

- Use a text editor to edit the **install-config.yaml** file prior to deploying your cluster and add one of the following stanzas:
 - To use shielded VMs for only control plane machines:

```
controlPlane:
  platform:
    gcp:
      secureBoot: Enabled
```

- To use shielded VMs for only compute machines:

```
compute:
- platform:
  gcp:
    secureBoot: Enabled
```

- To use shielded VMs for all machines:

```
platform:
  gcp:
    defaultMachinePlatform:
      secureBoot: Enabled
```

9.5.5.6. Enabling Confidential VMs

You can use Confidential VMs when installing your cluster. Confidential VMs encrypt data while it is being processed. For more information, see Google's documentation on [Confidential Computing](#). You can enable Confidential VMs and Shielded VMs at the same time, although they are not dependent on each other.



NOTE

Confidential VMs are currently not supported on 64-bit ARM architectures.

Prerequisites

- You have created an **install-config.yaml** file.

Procedure

- Use a text editor to edit the **install-config.yaml** file prior to deploying your cluster and add one of the following stanzas:
 - To use confidential VMs for only control plane machines:

```
controlPlane:
  platform:
```

```

gcp:
  confidentialCompute: Enabled ❶
  type: n2d-standard-8 ❷
  onHostMaintenance: Terminate ❸

```

- ❶ Enable confidential VMs.
- ❷ Specify a machine type that supports Confidential VMs. Confidential VMs require the N2D or C2D series of machine types. For more information on supported machine types, see [Supported operating systems and machine types](#).
- ❸ Specify the behavior of the VM during a host maintenance event, such as a hardware or software update. For a machine that uses Confidential VM, this value must be set to **Terminate**, which stops the VM. Confidential VMs do not support live VM migration.

- b. To use confidential VMs for only compute machines:

```

compute:
- platform:
  gcp:
    confidentialCompute: Enabled
    type: n2d-standard-8
    onHostMaintenance: Terminate

```

- c. To use confidential VMs for all machines:

```

platform:
  gcp:
    defaultMachinePlatform:
      confidentialCompute: Enabled
      type: n2d-standard-8
      onHostMaintenance: Terminate

```

9.5.5.7. Sample customized install-config.yaml file for GCP

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.



IMPORTANT

This sample YAML file is provided for reference only. You must obtain your **install-config.yaml** file by using the installation program and modify it.

```

apiVersion: v1
baseDomain: example.com ❶
credentialsMode: Mint ❷
controlPlane: ❸ ❹
  hyperthreading: Enabled ❺
  name: master
  platform:
    gcp:
      type: n2-standard-4

```

```
zones:
- us-central1-a
- us-central1-c
osDisk:
  diskType: pd-ssd
  diskSizeGB: 1024
  encryptionKey: 6
  kmsKey:
    name: worker-key
    keyRing: test-machine-keys
    location: global
    projectID: project-id
tags: 7
- control-plane-tag1
- control-plane-tag2
osImage: 8
  project: example-project-name
  name: example-image-name
replicas: 3
compute: 9 10
- hyperthreading: Enabled 11
  name: worker
  platform:
    gcp:
      type: n2-standard-4
      zones:
        - us-central1-a
        - us-central1-c
      osDisk:
        diskType: pd-standard
        diskSizeGB: 128
        encryptionKey: 12
        kmsKey:
          name: worker-key
          keyRing: test-machine-keys
          location: global
          projectID: project-id
      tags: 13
        - compute-tag1
        - compute-tag2
      osImage: 14
        project: example-project-name
        name: example-image-name
    replicas: 3
  metadata:
    name: test-cluster 15
  networking: 16
    clusterNetwork:
      - cidr: 10.128.0.0/14
      hostPrefix: 23
    machineNetwork:
      - cidr: 10.0.0.0/16
    networkType: OVNKubernetes 17
    serviceNetwork:
      - 172.30.0.0/16
```

```

platform:
  gcp:
    projectID: openshift-production 18
    region: us-central1 19
    defaultMachinePlatform:
      tags: 20
      - global-tag1
      - global-tag2
    osImage: 21
      project: example-project-name
      name: example-image-name
  pullSecret: '{"auths": ...}' 22
  fips: false 23
  sshKey: ssh-ed25519 AAAA... 24

```

- 1 15 18 19 22 Required. The installation program prompts you for this value.
- 2 Optional: Add this parameter to force the Cloud Credential Operator (CCO) to use the specified mode. By default, the CCO uses the root credentials in the **kube-system** namespace to dynamically try to determine the capabilities of the credentials. For details about CCO modes, see the "About the Cloud Credential Operator" section in the *Authentication and authorization* guide.
- 3 9 16 If you do not provide these parameters and values, the installation program provides the default value.
- 4 10 The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.
- 5 11 Whether to enable or disable simultaneous multithreading, or **hyperthreading**. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores. You can disable it by setting the parameter value to **Disabled**. If you disable simultaneous multithreading in some cluster machines, you must disable it in all cluster machines.



IMPORTANT

If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance. Use larger machine types, such as **n1-standard-8**, for your machines if you disable simultaneous multithreading.

- 6 12 Optional: The custom encryption key section to encrypt both virtual machines and persistent volumes. Your default compute service account must have the permissions granted to use your KMS key and have the correct IAM role assigned. The default service account name follows the **service-<project_number>@compute-system.iam.gserviceaccount.com** pattern. For more information about granting the correct permissions for your service account, see "Machine management" → "Creating compute machine sets" → "Creating a compute machine set on GCP".
- 7 13 20 Optional: A set of network tags to apply to the control plane or compute machine sets. The **platform.gcp.defaultMachinePlatform.tags** parameter will apply to both control plane and compute machines. If the **compute.platform.gcp.tags** or **controlPlane.platform.gcp.tags** parameters are set, they override the **platform.gcp.defaultMachinePlatform.tags** parameter.

- 8 14 21 Optional: A custom Red Hat Enterprise Linux CoreOS (RHCOS) that should be used to boot control plane and compute machines. The **project** and **name** parameters under
- 17 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 23 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.



IMPORTANT

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 24 You can optionally provide the **sshKey** value that you use to access the machines in your cluster.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

Additional resources

- [Enabling customer-managed encryption keys for a compute machine set](#)

9.5.5.8. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.



NOTE

The installation program does not support the proxy **readinessEndpoints** field.

**NOTE**

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.

**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

9.5.6. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.

**IMPORTANT**

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**. To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.



NOTE

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your **PATH**.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

9.5.7. Alternatives to storing administrator-level secrets in the kube-system project

By default, administrator secrets are stored in the **kube-system** project. If you configured the **credentialsMode** parameter in the **install-config.yaml** file to **Manual**, you must use one of the following alternatives:

- To manage long-term cloud credentials manually, follow the procedure in [Manually creating long-term credentials](#).
- To implement short-term credentials that are managed outside the cluster for individual components, follow the procedures in [Configuring a GCP cluster to use short-term credentials](#).

9.5.7.1. Manually creating long-term credentials

The Cloud Credential Operator (CCO) can be put into manual mode prior to installation in environments where the cloud identity and access management (IAM) APIs are not reachable, or the administrator prefers not to store an administrator-level credential secret in the cluster **kube-system** namespace.

Procedure

1. Add the following granular permissions to the GCP account that the installation program uses:

Example 9.30. Required GCP permissions

- compute.machineTypes.list
- compute.regions.list
- compute.zones.list
- dns.changes.create
- dns.changes.get
- dns.managedZones.create
- dns.managedZones.delete
- dns.managedZones.get
- dns.managedZones.list
- dns.networks.bindPrivateDNSZone
- dns.resourceRecordSets.create
- dns.resourceRecordSets.delete
- dns.resourceRecordSets.list

- If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```
apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...
```

- If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

- Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

- Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included 1 \
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml 2 \
  --to=<path_to_directory_for_credentials_requests> 3
```

- The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- Specify the location of the **install-config.yaml** file.
- Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

This command creates a YAML file for each **CredentialsRequest** object.

Sample CredentialsRequest object

```
apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
```

```

kind: GCPProviderSpec
predefinedRoles:
- roles/storage.admin
- roles/iam.serviceAccountUser
skipServiceCheck: true
...

```

6. Create YAML files for secrets in the **openshift-install** manifests directory that you generated previously. The secrets must be stored using the namespace and secret name defined in the **spec.secretRef** for each **CredentialsRequest** object.

Sample CredentialsRequest object with secrets

```

apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    ...
  secretRef:
    name: <component_secret>
    namespace: <component_namespace>
  ...

```

Sample Secret object

```

apiVersion: v1
kind: Secret
metadata:
  name: <component_secret>
  namespace: <component_namespace>
data:
  service_account.json: <base64_encoded_gcp_service_account_file>

```



IMPORTANT

Before upgrading a cluster that uses manually maintained credentials, you must ensure that the CCO is in an upgradeable state.

9.5.7.2. Configuring a GCP cluster to use short-term credentials

To install a cluster that is configured to use GCP Workload Identity, you must configure the CCO utility and create the required GCP resources for your cluster.

9.5.7.2.1. Configuring the Cloud Credential Operator utility

To create and manage cloud credentials from outside of the cluster when the Cloud Credential Operator (CCO) is operating in manual mode, extract and prepare the CCO utility (**ccoctl**) binary.

**NOTE**

The **ccoctl** utility is a Linux binary that must run in a Linux environment.

Prerequisites

- You have access to an OpenShift Container Platform account with cluster administrator access.
- You have installed the OpenShift CLI (**oc**).
- You have added one of the following authentication options to the GCP account that the installation program uses:
 - The **IAM Workload Identity Pool Admin** role.
 - The following granular permissions:

Example 9.31. Required GCP permissions

- `compute.projects.get`
- `iam.googleapis.com/workloadIdentityPoolProviders.create`
- `iam.googleapis.com/workloadIdentityPoolProviders.get`
- `iam.googleapis.com/workloadIdentityPools.create`
- `iam.googleapis.com/workloadIdentityPools.delete`
- `iam.googleapis.com/workloadIdentityPools.get`
- `iam.googleapis.com/workloadIdentityPools.undelete`
- `iam.roles.create`
- `iam.roles.delete`
- `iam.roles.list`
- `iam.roles.undelete`
- `iam.roles.update`
- `iam.serviceAccounts.create`
- `iam.serviceAccounts.delete`
- `iam.serviceAccounts.getIamPolicy`
- `iam.serviceAccounts.list`
- `iam.serviceAccounts.setIamPolicy`
- `iam.workloadIdentityPoolProviders.get`
- `iam.workloadIdentityPools.delete`
- `resourceManager.projects.get`

- resourcemanager.projects.getIamPolicy
- resourcemanager.projects.setIamPolicy
- storage.buckets.create
- storage.buckets.delete
- storage.buckets.get
- storage.buckets.getIamPolicy
- storage.buckets.setIamPolicy
- storage.objects.create
- storage.objects.delete
- storage.objects.list

Procedure

1. Set a variable for the OpenShift Container Platform release image by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

2. Obtain the CCO container image from the OpenShift Container Platform release image by running the following command:

```
$ CCO_IMAGE=$(oc adm release info --image-for='cloud-credential-operator'
$RELEASE_IMAGE -a ~/.pull-secret)
```



NOTE

Ensure that the architecture of the **\$RELEASE_IMAGE** matches the architecture of the environment in which you will use the **ccoctl** tool.

3. Extract the **ccoctl** binary from the CCO container image within the OpenShift Container Platform release image by running the following command:

```
$ oc image extract $CCO_IMAGE \
--file="/usr/bin/ccoctl.<rhel_version>" 1 \
-a ~/.pull-secret
```

- 1** For **<rhel_version>**, specify the value that corresponds to the version of Red Hat Enterprise Linux (RHEL) that the host uses. If no value is specified, **ccoctl.rhel8** is used by default. The following values are valid:

- **rhel8**: Specify this value for hosts that use RHEL 8.
- **rhel9**: Specify this value for hosts that use RHEL 9.

4. Change the permissions to make **ccoctl** executable by running the following command:

```
$ chmod 775 ccoctl.<rhel_version>
```

Verification

- To verify that **ccoctl** is ready to use, display the help file by running the following command:

```
$ ccoctl --help
```

Output of **ccoctl --help**

```
OpenShift credentials provisioning tool
```

```
Usage:
```

```
ccoctl [command]
```

```
Available Commands:
```

```
aws      Manage credentials objects for AWS cloud
azure    Manage credentials objects for Azure
gcp      Manage credentials objects for Google cloud
help     Help about any command
ibmcloud Manage credentials objects for IBM Cloud
nutanix  Manage credentials objects for Nutanix
```

```
Flags:
```

```
-h, --help  help for ccoctl
```

```
Use "ccoctl [command] --help" for more information about a command.
```

9.5.7.2.2. Creating GCP resources with the Cloud Credential Operator utility

You can use the **ccoctl gcp create-all** command to automate the creation of GCP resources.



NOTE

By default, **ccoctl** creates objects in the directory in which the commands are run. To create the objects in a different directory, use the **--output-dir** flag. This procedure uses **<path_to_ccoctl_output_dir>** to refer to this directory.

Prerequisites

You must have:

- Extracted and prepared the **ccoctl** binary.

Procedure

1. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

2. Extract the list of **CredentialsRequest** objects from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included \ 1
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \ 2
  --to=<path_to_directory_for_credentials_requests> \ 3
```

- 1 The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2 Specify the location of the **install-config.yaml** file.
- 3 Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.



NOTE

This command might take a few moments to run.

3. Use the **ccoctl** tool to process all **CredentialsRequest** objects by running the following command:

```
$ ccoctl gcp create-all \
  --name=<name> \ 1
  --region=<gcp_region> \ 2
  --project=<gcp_project_id> \ 3
  --credentials-requests-dir=<path_to_credentials_requests_directory> \ 4
```

- 1 Specify the user-defined name for all created GCP resources used for tracking.
- 2 Specify the GCP region in which cloud resources will be created.
- 3 Specify the GCP project ID in which cloud resources will be created.
- 4 Specify the directory containing the files of **CredentialsRequest** manifests to create GCP service accounts.



NOTE

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

Verification

- To verify that the OpenShift Container Platform secrets are created, list the files in the **<path_to_ccoctl_output_dir>/manifests** directory:

```
$ ls <path_to_ccoctl_output_dir>/manifests
```


Example output

```
cluster-authentication-02-config.yaml
openshift-cloud-controller-manager-gcp-ccm-cloud-credentials-credentials.yaml
openshift-cloud-credential-operator-cloud-credential-operator-gcp-ro-creds-credentials.yaml
openshift-cloud-network-config-controller-cloud-credentials-credentials.yaml
openshift-cluster-api-capg-manager-bootstrap-credentials-credentials.yaml
openshift-cluster-csi-drivers-gcp-pd-cloud-credentials-credentials.yaml
openshift-image-registry-installer-cloud-credentials-credentials.yaml
openshift-ingress-operator-cloud-credentials-credentials.yaml
openshift-machine-api-gcp-cloud-credentials-credentials.yaml
```

You can verify that the IAM service accounts are created by querying GCP. For more information, refer to GCP documentation on listing IAM service accounts.

9.5.7.2.3. Incorporating the Cloud Credential Operator utility manifests

To implement short-term security credentials managed outside the cluster for individual components, you must move the manifest files that the Cloud Credential Operator utility (**ccoctl**) created to the correct directories for the installation program.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have configured the Cloud Credential Operator utility (**ccoctl**).
- You have created the cloud provider resources that are required for your cluster with the **ccoctl** utility.

Procedure

1. Add the following granular permissions to the GCP account that the installation program uses:

Example 9.32. Required GCP permissions

- compute.machineTypes.list
- compute.regions.list
- compute.zones.list
- dns.changes.create
- dns.changes.get
- dns.managedZones.create
- dns.managedZones.delete
- dns.managedZones.get
- dns.managedZones.list
- dns.networks.bindPrivateDNSZone

- `dns.resourceRecordSets.create`
- `dns.resourceRecordSets.delete`
- `dns.resourceRecordSets.list`

- If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```
apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...
```

- If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

- Copy the manifests that the **ccoctl** utility generated to the **manifests** directory that the installation program created by running the following command:

```
$ cp /<path_to_ccoctl_output_dir>/manifests/* ./manifests/
```

- Copy the private key that the **ccoctl** utility generated in the **tls** directory to the installation directory by running the following command:

```
$ cp -a /<path_to_ccoctl_output_dir>/tls .
```

9.5.8. Network configuration phases

There are two phases prior to OpenShift Container Platform installation where you can customize the network configuration.

Phase 1

You can customize the following network-related fields in the **install-config.yaml** file before you create the manifest files:

- **networking.networkType**
- **networking.clusterNetwork**
- **networking.serviceNetwork**
- **networking.machineNetwork**

For more information on these fields, refer to *Installation configuration parameters*.

**NOTE**

Set the **networking.machineNetwork** to match the CIDR that the preferred NIC resides in.

**IMPORTANT**

The CIDR range **172.17.0.0/16** is reserved by libVirt. You cannot use this range or any range that overlaps with this range for any networks in your cluster.

Phase 2

After creating the manifest files by running **openshift-install create manifests**, you can define a customized Cluster Network Operator manifest with only the fields you want to modify. You can use the manifest to specify advanced network configuration.

You cannot override the values specified in phase 1 in the **install-config.yaml** file during phase 2. However, you can further customize the network plugin during phase 2.

9.5.9. Specifying advanced network configuration

You can use advanced network configuration for your network plugin to integrate your cluster into your existing network environment. You can specify advanced network configuration only before you install the cluster.

**IMPORTANT**

Customizing your network configuration by modifying the OpenShift Container Platform manifest files created by the installation program is not supported. Applying a manifest file that you create, as in the following procedure, is supported.

Prerequisites

- You have created the **install-config.yaml** file and completed any modifications to it.

Procedure

- Change to the directory that contains the installation program and create the manifests:

```
$ ./openshift-install create manifests --dir <installation_directory> 1
```

- 1** **<installation_directory>** specifies the name of the directory that contains the **install-config.yaml** file for your cluster.

- Create a stub manifest file for the advanced network configuration that is named **cluster-network-03-config.yml** in the **<installation_directory>/manifests/** directory:

```
apiVersion: operator.openshift.io/v1
kind: Network
metadata:
  name: cluster
spec:
```

- Specify the advanced network configuration for your cluster in the **cluster-network-03-config.yml** file, such as in the following example:

Enable IPsec for the OVN-Kubernetes network provider

```
apiVersion: operator.openshift.io/v1
kind: Network
metadata:
  name: cluster
spec:
  defaultNetwork:
    ovnKubernetesConfig:
      ipsecConfig:
        mode: Full
```

- Optional: Back up the **manifests/cluster-network-03-config.yml** file. The installation program consumes the **manifests/** directory when you create the Ignition config files.

9.5.10. Cluster Network Operator configuration

The configuration for the cluster network is specified as part of the Cluster Network Operator (CNO) configuration and stored in a custom resource (CR) object that is named **cluster**. The CR specifies the fields for the **Network** API in the **operator.openshift.io** API group.

The CNO configuration inherits the following fields during cluster installation from the **Network** API in the **Network.config.openshift.io** API group:

clusterNetwork

IP address pools from which pod IP addresses are allocated.

serviceNetwork

IP address pool for services.

defaultNetwork.type

Cluster network plugin. **OVNKubernetes** is the only supported plugin during installation.

You can specify the cluster network plugin configuration for your cluster by setting the fields for the **defaultNetwork** object in the CNO object named **cluster**.

9.5.10.1. Cluster Network Operator configuration object

The fields for the Cluster Network Operator (CNO) are described in the following table:

Table 9.7. Cluster Network Operator configuration object

Field	Type	Description
metadata.name	string	The name of the CNO object. This name is always cluster .


Field	Type	Description
spec.clusterNetwork	array	<p>A list specifying the blocks of IP addresses from which pod IP addresses are allocated and the subnet prefix length assigned to each individual node in the cluster. For example:</p> <pre>spec: clusterNetwork: - cidr: 10.128.0.0/19 hostPrefix: 23 - cidr: 10.128.32.0/19 hostPrefix: 23</pre>
spec.serviceNetwork	array	<p>A block of IP addresses for services. The OpenShift SDN and OVN-Kubernetes network plugins support only a single IP address block for the service network. For example:</p> <pre>spec: serviceNetwork: - 172.30.0.0/14</pre> <p>You can customize this field only in the install-config.yaml file before you create the manifests. The value is read-only in the manifest file.</p>
spec.defaultNetwork	object	Configures the network plugin for the cluster network.
spec.kubeProxy Config	object	The fields for this object specify the kube-proxy configuration. If you are using the OVN-Kubernetes cluster network plugin, the kube-proxy configuration has no effect.

defaultNetwork object configuration

The values for the **defaultNetwork** object are defined in the following table:

Table 9.8. **defaultNetwork** object

Field	Type	Description
-------	------	-------------

Field	Type	Description
type	string	<p>OVNKubernetes. The Red Hat OpenShift Networking network plugin is selected during installation. This value cannot be changed after cluster installation.</p>  <p>NOTE</p> <p>OpenShift Container Platform uses the OVN-Kubernetes network plugin by default. OpenShift SDN is no longer available as an installation choice for new clusters.</p>
ovnKubernetesConfig	object	This object is only valid for the OVN-Kubernetes network plugin.

Configuration for the OVN-Kubernetes network plugin

The following table describes the configuration fields for the OVN-Kubernetes network plugin:

Table 9.9. ovnKubernetesConfig object

Field	Type	Description
mtu	integer	<p>The maximum transmission unit (MTU) for the Geneve (Generic Network Virtualization Encapsulation) overlay network. This is detected automatically based on the MTU of the primary network interface. You do not normally need to override the detected MTU.</p> <p>If the auto-detected value is not what you expect it to be, confirm that the MTU on the primary network interface on your nodes is correct. You cannot use this option to change the MTU value of the primary network interface on the nodes.</p> <p>If your cluster requires different MTU values for different nodes, you must set this value to 100 less than the lowest MTU value in your cluster. For example, if some nodes in your cluster have an MTU of 9001, and some have an MTU of 1500, you must set this value to 1400.</p>
genevePort	integer	The port to use for all Geneve packets. The default value is 6081 . This value cannot be changed after cluster installation.
ipsecConfig	object	Specify a configuration object for customizing the IPsec configuration.


Field	Type	Description
ipv4	object	Specifies a configuration object for IPv4 settings.
ipv6	object	Specifies a configuration object for IPv6 settings.
policyAuditConfig	object	Specify a configuration object for customizing network policy audit logging. If unset, the defaults audit log settings are used.
gatewayConfig	object	Optional: Specify a configuration object for customizing how egress traffic is sent to the node gateway.  <p>NOTE</p> <p>While migrating egress traffic, you can expect some disruption to workloads and service traffic until the Cluster Network Operator (CNO) successfully rolls out the changes.</p>

Table 9.10. `ovnKubernetesConfig.ipv4` object

Field	Type	Description
internalTransitSwitchSubnet	string	If your existing network infrastructure overlaps with the 100.88.0.0/16 IPv4 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. The subnet for the distributed transit switch that enables east-west traffic. This subnet cannot overlap with any other subnets used by OVN-Kubernetes or on the host itself. It must be large enough to accommodate one IP address per node in your cluster. The default value is 100.88.0.0/16 .
internalJoinSubnet	string	If your existing network infrastructure overlaps with the 100.64.0.0/16 IPv4 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. You must ensure that the IP address range does not overlap with any other subnet used by your OpenShift Container Platform installation. The IP address range must be larger than the maximum number of nodes that can be added to the cluster. For example, if the clusterNetwork.cidr value is 10.128.0.0/14 and the clusterNetwork.hostPrefix value is /23 , then the maximum number of nodes is 2⁽²³⁻¹⁴⁾=512 . The default value is 100.64.0.0/16 .

Table 9.11. `ovnKubernetesConfig.ipv6` object

Field	Type	Description
internalTransitSwitchSubnet	string	<p>If your existing network infrastructure overlaps with the fd97::/64 IPv6 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. The subnet for the distributed transit switch that enables east-west traffic. This subnet cannot overlap with any other subnets used by OVN-Kubernetes or on the host itself. It must be large enough to accommodate one IP address per node in your cluster.</p> <p>The default value is fd97::/64.</p>
internalJoinSubnet	string	<p>If your existing network infrastructure overlaps with the fd98::/64 IPv6 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. You must ensure that the IP address range does not overlap with any other subnet used by your OpenShift Container Platform installation. The IP address range must be larger than the maximum number of nodes that can be added to the cluster.</p> <p>The default value is fd98::/64.</p>

Table 9.12. **policyAuditConfig** object

Field	Type	Description
rateLimit	integer	The maximum number of messages to generate every second per node. The default value is 20 messages per second.
maxFileSize	integer	The maximum size for the audit log in bytes. The default value is 50000000 or 50 MB.
maxLogFiles	integer	The maximum number of log files that are retained.
destination	string	<p>One of the following additional audit log targets:</p> <p>libc The libc syslog() function of the journald process on the host.</p> <p>udp:<host>:<port> A syslog server. Replace <host>:<port> with the host and port of the syslog server.</p> <p>unix:<file> A Unix Domain Socket file specified by <file>.</p> <p>null Do not send the audit logs to any additional target.</p>
syslogFacility	string	The syslog facility, such as kern , as defined by RFC5424. The default value is local0 .

Table 9.13. gatewayConfig object

Field	Type	Description
routingViaHost	boolean	<p>Set this field to true to send egress traffic from pods to the host networking stack. For highly-specialized installations and applications that rely on manually configured routes in the kernel routing table, you might want to route egress traffic to the host networking stack. By default, egress traffic is processed in OVN to exit the cluster and is not affected by specialized routes in the kernel routing table. The default value is false.</p> <p>This field has an interaction with the Open vSwitch hardware offloading feature. If you set this field to true, you do not receive the performance benefits of the offloading because egress traffic is processed by the host networking stack.</p>
ipForwarding	object	<p>You can control IP forwarding for all traffic on OVN-Kubernetes managed interfaces by using the ipForwarding specification in the Network resource. Specify Restricted to only allow IP forwarding for Kubernetes related traffic. Specify Global to allow forwarding of all IP traffic. For new installations, the default is Restricted. For updates to OpenShift Container Platform 4.14 or later, the default is Global.</p>
ipv4	object	<p>Optional: Specify an object to configure the internal OVN-Kubernetes masquerade address for host to service traffic for IPv4 addresses.</p>
ipv6	object	<p>Optional: Specify an object to configure the internal OVN-Kubernetes masquerade address for host to service traffic for IPv6 addresses.</p>

Table 9.14. gatewayConfig.ipv4 object

Field	Type	Description
internalMasqueradeSubnet	string	<p>The masquerade IPv4 addresses that are used internally to enable host to service traffic. The host is configured with these IP addresses as well as the shared gateway bridge interface. The default value is 169.254.169.0/29.</p>

Table 9.15. gatewayConfig.ipv6 object

Field	Type	Description
internalMasqueradeSubnet	string	<p>The masquerade IPv6 addresses that are used internally to enable host to service traffic. The host is configured with these IP addresses as well as the shared gateway bridge interface. The default value is fd69::/125.</p>

Table 9.16. `ipsecConfig` object

Field	Type	Description
<code>mode</code>	<code>string</code>	<p>Specifies the behavior of the IPsec implementation. Must be one of the following values:</p> <ul style="list-style-type: none"> ● Disabled: IPsec is not enabled on cluster nodes. ● External: IPsec is enabled for network traffic with external hosts. ● Full: IPsec is enabled for pod traffic and network traffic with external hosts.

Example OVN-Kubernetes configuration with IPsec enabled

```

defaultNetwork:
  type: OVNKubernetes
  ovnKubernetesConfig:
    mtu: 1400
    genevePort: 6081
    ipsecConfig:
      mode: Full

```




IMPORTANT

Using OVNKubernetes can lead to a stack exhaustion problem on IBM Power®.

kubeProxyConfig object configuration (OpenShiftSDN container network interface only)

The values for the `kubeProxyConfig` object are defined in the following table:


Table 9.17. `kubeProxyConfig` object

Field	Type	Description
<code>iptablesSyncPeriod</code>	<code>string</code>	<p>The refresh period for <code>iptables</code> rules. The default value is 30s. Valid suffixes include s, m, and h and are described in the Go time package documentation.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p>NOTE</p> <p>Because of performance improvements introduced in OpenShift Container Platform 4.3 and greater, adjusting the <code>iptablesSyncPeriod</code> parameter is no longer necessary.</p> </div> </div>

Field	Type	Description
proxyArguments.iptables-min-sync-period	array	<p>The minimum duration before refreshing iptables rules. This field ensures that the refresh does not happen too frequently. Valid suffixes include s, m, and h and are described in the Go time package. The default value is:</p> <pre>kubeProxyConfig: proxyArguments: iptables-min-sync-period: - 0s</pre>

9.5.11. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.

Procedure

1. Remove any existing GCP credentials that do not use the service account key for the GCP account that you configured for your cluster and that are stored in the following locations:
 - The **GOOGLE_CREDENTIALS**, **GOOGLE_CLOUD_KEYFILE_JSON**, or **GCPLOUD_KEYFILE_JSON** environment variables
 - The `~/gcp/osServiceAccount.json` file
 - The **gcloud cli** default credentials
2. Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

- 1 For `<installation_directory>`, specify the location of your customized `./install-config.yaml` file.
 - 2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.
3. Optional: You can reduce the number of permissions for the service account that you used to install the cluster.
- If you assigned the **Owner** role to your service account, you can remove that role and replace it with the **Viewer** role.
 - If you included the **Service Account Key Admin** role, you can remove it.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to `<installation_directory>/openshift_install.log`.



IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

9.5.12. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

Additional resources

- See [Accessing the web console](#) for more details about accessing and understanding the OpenShift Container Platform web console.

9.5.13. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service

9.5.14. Next steps

- [Customize your cluster](#).

- If necessary, you can [opt out of remote health reporting](#) .

9.6. INSTALLING A CLUSTER ON GCP IN A RESTRICTED NETWORK

In OpenShift Container Platform 4.16, you can install a cluster on Google Cloud Platform (GCP) in a restricted network by creating an internal mirror of the installation release content on an existing Google Virtual Private Cloud (VPC).




IMPORTANT

You can install an OpenShift Container Platform cluster by using mirrored installation release content, but your cluster will require internet access to use the GCP APIs.

9.6.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [configured a GCP project](#) to host the cluster.
- You [mirrored the images for a disconnected installation](#) to your registry and obtained the **imageContentSources** data for your version of OpenShift Container Platform.



IMPORTANT

Because the installation media is on the mirror host, you can use that computer to complete all installation steps.

- You have an existing VPC in GCP. While installing a cluster in a restricted network that uses installer-provisioned infrastructure, you cannot use the installer-provisioned VPC. You must use a user-provisioned VPC that satisfies one of the following requirements:
 - Contains the mirror registry
 - Has firewall rules or a peering connection to access the mirror registry hosted elsewhere
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to. While you might need to grant access to more sites, you must grant access to ***.googleapis.com** and **accounts.google.com**.

9.6.2. About installations in restricted networks

In OpenShift Container Platform 4.16, you can perform an installation that does not require an active connection to the internet to obtain software components. Restricted network installations can be completed using installer-provisioned infrastructure or user-provisioned infrastructure, depending on the cloud platform to which you are installing the cluster.

If you choose to perform a restricted network installation on a cloud platform, you still require access to its cloud APIs. Some cloud functions, like Amazon Web Service's Route 53 DNS and IAM services, require internet access. Depending on your network, you might require less internet access for an installation on bare metal hardware, Nutanix, or on VMware vSphere.

To complete a restricted network installation, you must create a registry that mirrors the contents of the OpenShift image registry and contains the installation media. You can create this registry on a mirror host, which can access both the internet and your closed network, or by using other methods that meet your restrictions.

9.6.2.1. Additional limits

Clusters in restricted networks have the following additional limitations and restrictions:

- The **ClusterVersion** status includes an **Unable to retrieve available updates** error.
- By default, you cannot use the contents of the Developer Catalog because you cannot access the required image stream tags.

9.6.3. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to obtain the images that are necessary to install your cluster.

You must have internet access to:


- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.

9.6.4. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.



NOTE

You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.



NOTE

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.



NOTE

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```



NOTE

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:


```
$ ssh-add <path>/<file_name> 1
```

- 1 Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

9.6.5. Creating the installation configuration file

You can customize the OpenShift Container Platform cluster you install on Google Cloud Platform (GCP).

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster. For a restricted network installation, these files are on your mirror host.
- You have the **imageContentSources** values that were generated during mirror registry creation.
- You have obtained the contents of the certificate for your mirror registry.

Procedure

1. Create the **install-config.yaml** file.
 - a. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install create install-config --dir <installation_directory> 1
```

- 1 For **<installation_directory>**, specify the directory name to store the files that the installation program creates.

When specifying the directory:

- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
- Use an empty directory. Some installation assets, such as bootstrap X.509 certificates, have short expiration intervals, therefore you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

- b. At the prompts, provide the configuration details for your cloud:
 - i. Optional: Select an SSH key to use to access your cluster machines.

**NOTE**

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- ii. Select **gcp** as the platform to target.
 - iii. If you have not configured the service account key for your GCP account on your computer, you must obtain it from GCP and paste the contents of the file or enter the absolute path to the file.
 - iv. Select the project ID to provision the cluster in. The default value is specified by the service account that you configured.
 - v. Select the region to deploy the cluster to.
 - vi. Select the base domain to deploy the cluster to. The base domain corresponds to the public DNS zone that you created for your cluster.
 - vii. Enter a descriptive name for your cluster.
2. Edit the **install-config.yaml** file to give the additional information that is required for an installation in a restricted network.
 - a. Update the **pullSecret** value to contain the authentication information for your registry:

```
pullSecret: '{"auths":{"<mirror_host_name>:5000":{"auth": "<credentials>","email":
"you@example.com"}}}'
```

For **<mirror_host_name>**, specify the registry domain name that you specified in the certificate for your mirror registry, and for **<credentials>**, specify the base64-encoded user name and password for your mirror registry.

- b. Add the **additionalTrustBundle** parameter and value.

```
additionalTrustBundle: |
  -----BEGIN CERTIFICATE-----
  /-----END CERTIFICATE-----
```

The value must be the contents of the certificate file that you used for your mirror registry. The certificate file can be an existing, trusted certificate authority, or the self-signed certificate that you generated for the mirror registry.

- c. Define the network and subnets for the VPC to install the cluster in under the parent **platform.gcp** field:

```
network: <existing_vpc>
controlPlaneSubnet: <control_plane_subnet>
computeSubnet: <compute_subnet>
```

■

For **platform.gcp.network**, specify the name for the existing Google VPC. For **platform.gcp.controlPlaneSubnet** and **platform.gcp.computeSubnet**, specify the existing subnets to deploy the control plane machines and compute machines, respectively.

- d. Add the image content resources, which resemble the following YAML excerpt:

```
imageContentSources:
- mirrors:
  - <mirror_host_name>:5000/<repo_name>/release
    source: quay.io/openshift-release-dev/ocp-release
- mirrors:
  - <mirror_host_name>:5000/<repo_name>/release
    source: registry.redhat.io/ocp/release
```


For these values, use the **imageContentSources** that you recorded during mirror registry creation.

- e. Optional: Set the publishing strategy to **Internal**:

```
publish: Internal
```

By setting this option, you create an internal Ingress Controller and a private load balancer.

3. Make any other modifications to the **install-config.yaml** file that you require. For more information about the parameters, see "Installation configuration parameters".
4. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

Additional resources

- [Installation configuration parameters for GCP](#)

9.6.5.1. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 9.18. Minimum resource requirements

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Bootstrap	RHCOS	4	16 GB	100 GB	300
Control plane	RHCOS	4	16 GB	100 GB	300

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Compute	RHCOS, RHEL 8.6 and later [3]	2	8 GB	100 GB	300

1. One vCPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio: $(\text{threads per core} \times \text{cores}) \times \text{sockets} = \text{vCPUs}$.
2. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes which require a 10 ms p99 fsync duration. Note that on many cloud platforms, storage size and IOPS scale together, so you might need to over-allocate storage volume to obtain sufficient performance.
3. As with all user-provisioned installations, if you choose to use RHEL compute machines in your cluster, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Use of RHEL 7 compute machines is deprecated and has been removed in OpenShift Container Platform 4.10 and later.

NOTE

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

9.6.5.2. Tested instance types for GCP

The following Google Cloud Platform instance types have been tested with OpenShift Container Platform.

Example 9.33. Machine series

- **A2**
- **A3**
- **C2**
- **C2D**
- **C3**
- **C3D**
- **E2**
- **M1**
- **N1**
- **N2**
- **N2D**
- **Tau T2D**

9.6.5.3. Tested instance types for GCP on 64-bit ARM infrastructures

The following Google Cloud Platform (GCP) 64-bit ARM instance types have been tested with OpenShift Container Platform.

Example 9.34. Machine series for 64-bit ARM machines

- **Tau T2A**

9.6.5.4. Using custom machine types

Using a custom machine type to install a OpenShift Container Platform cluster is supported.

Consider the following when using a custom machine type:

- Similar to predefined instance types, custom machine types must meet the minimum resource requirements for control plane and compute machines. For more information, see "Minimum resource requirements for cluster installation".
- The name of the custom machine type must adhere to the following syntax:
custom-`<number_of_cpus>`-`<amount_of_memory_in_mb>`

For example, **custom-6-20480**.

As part of the installation process, you specify the custom machine type in the **install-config.yaml** file.

Sample install-config.yaml file with a custom machine type

```
compute:
  - architecture: amd64
```

```

hyperthreading: Enabled
name: worker
platform:
  gcp:
    type: custom-6-20480
replicas: 2
controlPlane:
  architecture: amd64
  hyperthreading: Enabled
  name: master
  platform:
    gcp:
      type: custom-6-20480
  replicas: 3

```

9.6.5.5. Enabling Shielded VMs

You can use Shielded VMs when installing your cluster. Shielded VMs have extra security features including secure boot, firmware and integrity monitoring, and rootkit detection. For more information, see Google's documentation on [Shielded VMs](#).



NOTE

Shielded VMs are currently not supported on clusters with 64-bit ARM infrastructures.

Prerequisites

- You have created an **install-config.yaml** file.

Procedure

- Use a text editor to edit the **install-config.yaml** file prior to deploying your cluster and add one of the following stanzas:
 - To use shielded VMs for only control plane machines:

```

controlPlane:
  platform:
    gcp:
      secureBoot: Enabled

```

- To use shielded VMs for only compute machines:

```

compute:
  - platform:
      gcp:
        secureBoot: Enabled

```

- To use shielded VMs for all machines:


```

platform:
  gcp:
    defaultMachinePlatform:
      secureBoot: Enabled

```

9.6.5.6. Enabling Confidential VMs

You can use Confidential VMs when installing your cluster. Confidential VMs encrypt data while it is being processed. For more information, see Google's documentation on [Confidential Computing](#). You can enable Confidential VMs and Shielded VMs at the same time, although they are not dependent on each other.



NOTE

Confidential VMs are currently not supported on 64-bit ARM architectures.

Prerequisites

- You have created an **install-config.yaml** file.

Procedure

- Use a text editor to edit the **install-config.yaml** file prior to deploying your cluster and add one of the following stanzas:

- To use confidential VMs for only control plane machines:

```
controlPlane:
  platform:
    gcp:
      confidentialCompute: Enabled 1
      type: n2d-standard-8 2
      onHostMaintenance: Terminate 3
```

- Enable confidential VMs.
- Specify a machine type that supports Confidential VMs. Confidential VMs require the N2D or C2D series of machine types. For more information on supported machine types, see [Supported operating systems and machine types](#).
- Specify the behavior of the VM during a host maintenance event, such as a hardware or software update. For a machine that uses Confidential VM, this value must be set to **Terminate**, which stops the VM. Confidential VMs do not support live VM migration.

- To use confidential VMs for only compute machines:

```
compute:
- platform:
  gcp:
    confidentialCompute: Enabled
    type: n2d-standard-8
    onHostMaintenance: Terminate
```

- To use confidential VMs for all machines:

```
platform:
  gcp:
    defaultMachinePlatform:
```

```
confidentialCompute: Enabled
type: n2d-standard-8
onHostMaintenance: Terminate
```

9.6.5.7. Sample customized install-config.yaml file for GCP

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.



IMPORTANT

This sample YAML file is provided for reference only. You must obtain your **install-config.yaml** file by using the installation program and modify it.

```
apiVersion: v1
baseDomain: example.com 1
credentialsMode: Mint 2
controlPlane: 3 4
  hyperthreading: Enabled 5
  name: master
  platform:
    gcp:
      type: n2-standard-4
      zones:
      - us-central1-a
      - us-central1-c
    osDisk:
      diskType: pd-ssd
      diskSizeGB: 1024
      encryptionKey: 6
      kmsKey:
        name: worker-key
        keyRing: test-machine-keys
        location: global
        projectID: project-id
    tags: 7
    - control-plane-tag1
    - control-plane-tag2
    osImage: 8
      project: example-project-name
      name: example-image-name
  replicas: 3
compute: 9 10
- hyperthreading: Enabled 11
  name: worker
  platform:
    gcp:
      type: n2-standard-4
      zones:
      - us-central1-a
      - us-central1-c
    osDisk:
      diskType: pd-standard
```



```

diskSizeGB: 128
encryptionKey: 12
  kmsKey:
    name: worker-key
    keyRing: test-machine-keys
    location: global
    projectID: project-id
tags: 13
- compute-tag1
- compute-tag2
osImage: 14
  project: example-project-name
  name: example-image-name
replicas: 3
metadata:
  name: test-cluster 15
networking:
  clusterNetwork:
    - cidr: 10.128.0.0/14
    hostPrefix: 23
  machineNetwork:
    - cidr: 10.0.0.0/16
  networkType: OVNKubernetes 16
  serviceNetwork:
    - 172.30.0.0/16
platform:
  gcp:
    projectID: openshift-production 17
    region: us-central1 18
  defaultMachinePlatform:
    tags: 19
    - global-tag1
    - global-tag2
    osImage: 20
    project: example-project-name
    name: example-image-name
  network: existing_vpc 21
  controlPlaneSubnet: control_plane_subnet 22
  computeSubnet: compute_subnet 23
pullSecret: '{"auths":{"<local_registry>":{"auth": "<credentials>","email": "you@example.com"}}}' 24
fips: false 25
sshKey: ssh-ed25519 AAAA... 26
additionalTrustBundle: | 27
  -----BEGIN CERTIFICATE-----
  <MY_TRUSTED_CA_CERT>
  -----END CERTIFICATE-----
imageContentSources: 28
- mirrors:
  - <local_registry>/<local_repository_name>/release
  source: quay.io/openshift-release-dev/ocp-release
- mirrors:
  - <local_registry>/<local_repository_name>/release
  source: quay.io/openshift-release-dev/ocp-v4.0-art-dev

```

- 1 15 17 18 Required. The installation program prompts you for this value.
- 2 Optional: Add this parameter to force the Cloud Credential Operator (CCO) to use the specified mode. By default, the CCO uses the root credentials in the **kube-system** namespace to dynamically try to determine the capabilities of the credentials. For details about CCO modes, see the "About the Cloud Credential Operator" section in the *Authentication and authorization* guide.
- 3 9 If you do not provide these parameters and values, the installation program provides the default value.
- 4 10 The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.
- 5 11 Whether to enable or disable simultaneous multithreading, or **hyperthreading**. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores. You can disable it by setting the parameter value to **Disabled**. If you disable simultaneous multithreading in some cluster machines, you must disable it in all cluster machines.



IMPORTANT

If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance. Use larger machine types, such as **n1-standard-8**, for your machines if you disable simultaneous multithreading.

- 6 12 Optional: The custom encryption key section to encrypt both virtual machines and persistent volumes. Your default compute service account must have the permissions granted to use your KMS key and have the correct IAM role assigned. The default service account name follows the **service-<project_number>@compute-system.iam.gserviceaccount.com** pattern. For more information about granting the correct permissions for your service account, see "Machine management" → "Creating compute machine sets" → "Creating a compute machine set on GCP".
- 7 13 19 Optional: A set of network tags to apply to the control plane or compute machine sets. The **platform.gcp.defaultMachinePlatform.tags** parameter will apply to both control plane and compute machines. If the **compute.platform.gcp.tags** or **controlPlane.platform.gcp.tags** parameters are set, they override the **platform.gcp.defaultMachinePlatform.tags** parameter.
- 8 14 20 Optional: A custom Red Hat Enterprise Linux CoreOS (RHCOS) that should be used to boot control plane and compute machines. The **project** and **name** parameters under **platform.gcp.defaultMachinePlatform.osImage** apply to both control plane and compute machines. If the **project** and **name** parameters under **controlPlane.platform.gcp.osImage** or **compute.platform.gcp.osImage** are set, they override the **platform.gcp.defaultMachinePlatform.osImage** parameters.
- 16 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 21 Specify the name of an existing VPC.
- 22 Specify the name of the existing subnet to deploy the control plane machines to. The subnet must belong to the VPC that you specified.
- 23 Specify the name of the existing subnet to deploy the compute machines to. The subnet must belong to the VPC that you specified.


- 24 For **<local_registry>**, specify the registry domain name, and optionally the port, that your mirror registry uses to serve content. For example, **registry.example.com** or
- 25 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.



IMPORTANT

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 26 You can optionally provide the **sshKey** value that you use to access the machines in your cluster.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- 27 Provide the contents of the certificate file that you used for your mirror registry.
- 28 Provide the **imageContentSources** section from the output of the command to mirror the repository.

9.6.5.8. Create an Ingress Controller with global access on GCP

You can create an Ingress Controller that has global access to a Google Cloud Platform (GCP) cluster. Global access is only available to Ingress Controllers using internal load balancers.

Prerequisites

- You created the **install-config.yaml** and complete any modifications to it.

Procedure

Create an Ingress Controller with global access on a new GCP cluster.

- Change to the directory that contains the installation program and create a manifest file:

```
$ ./openshift-install create manifests --dir <installation_directory> 1
```

- For **<installation_directory>**, specify the name of the directory that contains the **install-config.yaml** file for your cluster.
- Create a file that is named **cluster-ingress-default-ingresscontroller.yaml** in the **<installation_directory>/manifests/** directory:

```
$ touch <installation_directory>/manifests/cluster-ingress-default-ingresscontroller.yaml 1
```

- 1** For **<installation_directory>**, specify the directory name that contains the **manifests/** directory for your cluster.

After creating the file, several network configuration files are in the **manifests/** directory, as shown:

```
$ ls <installation_directory>/manifests/cluster-ingress-default-ingresscontroller.yaml
```

Example output

```
cluster-ingress-default-ingresscontroller.yaml
```

- Open the **cluster-ingress-default-ingresscontroller.yaml** file in an editor and enter a custom resource (CR) that describes the Operator configuration you want:

Sample `clientAccess` configuration to Global

```
apiVersion: operator.openshift.io/v1
kind: IngressController
metadata:
  name: default
  namespace: openshift-ingress-operator
spec:
  endpointPublishingStrategy:
    loadBalancer:
      providerParameters:
        gcp:
          clientAccess: Global 1
          type: GCP
          scope: Internal 2
          type: LoadBalancerService
```

- 1** Set **gcp.clientAccess** to **Global**.
- 2** Global access is only available to Ingress Controllers using internal load balancers.

9.6.5.9. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to

bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.

**NOTE**

The installation program does not support the proxy **readinessEndpoints** field.

**NOTE**

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
┌ $ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.

**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

9.6.6. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.

**IMPORTANT**

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
┌ $ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**.
To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:


```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.



NOTE

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your **PATH**.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

9.6.7. Alternatives to storing administrator-level secrets in the kube-system project

By default, administrator secrets are stored in the **kube-system** project. If you configured the **credentialsMode** parameter in the **install-config.yaml** file to **Manual**, you must use one of the following alternatives:

- To manage long-term cloud credentials manually, follow the procedure in [Manually creating long-term credentials](#).
- To implement short-term credentials that are managed outside the cluster for individual components, follow the procedures in [Configuring a GCP cluster to use short-term credentials](#).

9.6.7.1. Manually creating long-term credentials

The Cloud Credential Operator (CCO) can be put into manual mode prior to installation in environments where the cloud identity and access management (IAM) APIs are not reachable, or the administrator prefers not to store an administrator-level credential secret in the cluster **kube-system** namespace.

Procedure

1. Add the following granular permissions to the GCP account that the installation program uses:

Example 9.35. Required GCP permissions

- compute.machineTypes.list
- compute.regions.list
- compute.zones.list
- dns.changes.create
- dns.changes.get
- dns.managedZones.create
- dns.managedZones.delete
- dns.managedZones.get
- dns.managedZones.list
- dns.networks.bindPrivateDNSZone
- dns.resourceRecordSets.create
- dns.resourceRecordSets.delete

- `dns.resourceRecordSets.list`

- If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```
apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...
```

- If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

- Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

- Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included 1 \
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml 2 \
  --to=<path_to_directory_for_credentials_requests> 3
```

- 1** The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2** Specify the location of the **install-config.yaml** file.
- 3** Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

This command creates a YAML file for each **CredentialsRequest** object.

Sample CredentialsRequest object

```
apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
```

```
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: GCPProviderSpec
    predefinedRoles:
      - roles/storage.admin
      - roles/iam.serviceAccountUser
    skipServiceCheck: true
  ...
```


6. Create YAML files for secrets in the **openshift-install** manifests directory that you generated previously. The secrets must be stored using the namespace and secret name defined in the **spec.secretRef** for each **CredentialsRequest** object.

Sample CredentialsRequest object with secrets

```
apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    ...
  secretRef:
    name: <component_secret>
    namespace: <component_namespace>
  ...
```

Sample Secret object

```
apiVersion: v1
kind: Secret
metadata:
  name: <component_secret>
  namespace: <component_namespace>
data:
  service_account.json: <base64_encoded_gcp_service_account_file>
```



IMPORTANT

Before upgrading a cluster that uses manually maintained credentials, you must ensure that the CCO is in an upgradeable state.

9.6.7.2. Configuring a GCP cluster to use short-term credentials

To install a cluster that is configured to use GCP Workload Identity, you must configure the CCO utility and create the required GCP resources for your cluster.

9.6.7.2.1. Configuring the Cloud Credential Operator utility

To create and manage cloud credentials from outside of the cluster when the Cloud Credential Operator (CCO) is operating in manual mode, extract and prepare the CCO utility (**ccoctl**) binary.



NOTE

The **ccoctl** utility is a Linux binary that must run in a Linux environment.

Prerequisites

- You have access to an OpenShift Container Platform account with cluster administrator access.
- You have installed the OpenShift CLI (**oc**).
- You have added one of the following authentication options to the GCP account that the installation program uses:
 - The **IAM Workload Identity Pool Admin** role.
 - The following granular permissions:

Example 9.36. Required GCP permissions

- `compute.projects.get`
- `iam.googleapis.com/workloadIdentityPoolProviders.create`
- `iam.googleapis.com/workloadIdentityPoolProviders.get`
- `iam.googleapis.com/workloadIdentityPools.create`
- `iam.googleapis.com/workloadIdentityPools.delete`
- `iam.googleapis.com/workloadIdentityPools.get`
- `iam.googleapis.com/workloadIdentityPools.undelete`
- `iam.roles.create`
- `iam.roles.delete`
- `iam.roles.list`
- `iam.roles.undelete`
- `iam.roles.update`
- `iam.serviceAccounts.create`
- `iam.serviceAccounts.delete`
- `iam.serviceAccounts.getIamPolicy`
- `iam.serviceAccounts.list`
- `iam.serviceAccounts.setIamPolicy`
- `iam.workloadIdentityPoolProviders.get`

- iam.workloadIdentityPools.delete
- resourcemanager.projects.get
- resourcemanager.projects.getIamPolicy
- resourcemanager.projects.setIamPolicy
- storage.buckets.create
- storage.buckets.delete
- storage.buckets.get
- storage.buckets.getIamPolicy
- storage.buckets.setIamPolicy
- storage.objects.create
- storage.objects.delete
- storage.objects.list

Procedure

1. Set a variable for the OpenShift Container Platform release image by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

2. Obtain the CCO container image from the OpenShift Container Platform release image by running the following command:

```
$ CCO_IMAGE=$(oc adm release info --image-for='cloud-credential-operator'
$RELEASE_IMAGE -a ~/.pull-secret)
```



NOTE

Ensure that the architecture of the **\$RELEASE_IMAGE** matches the architecture of the environment in which you will use the **ccoctl** tool.

3. Extract the **ccoctl** binary from the CCO container image within the OpenShift Container Platform release image by running the following command:

```
$ oc image extract $CCO_IMAGE \
--file="/usr/bin/ccoctl.<rhel_version>" \
-a ~/.pull-secret
```

- 1 For **<rhel_version>**, specify the value that corresponds to the version of Red Hat Enterprise Linux (RHEL) that the host uses. If no value is specified, **ccoctl.rhel8** is used by default. The following values are valid:

- **rhel8**: Specify this value for hosts that use RHEL 8.
- **rhel9**: Specify this value for hosts that use RHEL 9.

4. Change the permissions to make **ccoctl** executable by running the following command:

```
$ chmod 775 ccoctl.<rhel_version>
```

Verification

- To verify that **ccoctl** is ready to use, display the help file by running the following command:

```
$ ccoctl --help
```

Output of **ccoctl --help**

```
OpenShift credentials provisioning tool
```

```
Usage:
```

```
ccoctl [command]
```

```
Available Commands:
```

```
aws      Manage credentials objects for AWS cloud
azure    Manage credentials objects for Azure
gcp      Manage credentials objects for Google cloud
help     Help about any command
ibmcloud Manage credentials objects for IBM Cloud
nutanix  Manage credentials objects for Nutanix
```


```
Flags:
```

```
-h, --help  help for ccoctl
```

```
Use "ccoctl [command] --help" for more information about a command.
```

9.6.7.2.2. Creating GCP resources with the Cloud Credential Operator utility

You can use the **ccoctl gcp create-all** command to automate the creation of GCP resources.



NOTE

By default, **ccoctl** creates objects in the directory in which the commands are run. To create the objects in a different directory, use the **--output-dir** flag. This procedure uses **<path_to_ccoctl_output_dir>** to refer to this directory.

Prerequisites

You must have:

- Extracted and prepared the **ccoctl** binary.

Procedure


1. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

2. Extract the list of **CredentialsRequest** objects from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included \1
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \2
  --to=<path_to_directory_for_credentials_requests> \3
```

- 1 The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2 Specify the location of the **install-config.yaml** file.
- 3 Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.



NOTE

This command might take a few moments to run.

3. Use the **ccoctl** tool to process all **CredentialsRequest** objects by running the following command:

```
$ ccoctl gcp create-all \
  --name=<name> \1
  --region=<gcp_region> \2
  --project=<gcp_project_id> \3
  --credentials-requests-dir=<path_to_credentials_requests_directory> \4
```

- 1 Specify the user-defined name for all created GCP resources used for tracking.
- 2 Specify the GCP region in which cloud resources will be created.
- 3 Specify the GCP project ID in which cloud resources will be created.
- 4 Specify the directory containing the files of **CredentialsRequest** manifests to create GCP service accounts.



NOTE

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

Verification

- To verify that the OpenShift Container Platform secrets are created, list the files in the `<path_to_ccoctl_output_dir>/manifests` directory:

```
$ ls <path_to_ccoctl_output_dir>/manifests
```

Example output

```
cluster-authentication-02-config.yaml
openshift-cloud-controller-manager-gcp-ccm-cloud-credentials-credentials.yaml
openshift-cloud-credential-operator-cloud-credential-operator-gcp-ro-creds-credentials.yaml
openshift-cloud-network-config-controller-cloud-credentials-credentials.yaml
openshift-cluster-api-capg-manager-bootstrap-credentials-credentials.yaml
openshift-cluster-csi-drivers-gcp-pd-cloud-credentials-credentials.yaml
openshift-image-registry-installer-cloud-credentials-credentials.yaml
openshift-ingress-operator-cloud-credentials-credentials.yaml
openshift-machine-api-gcp-cloud-credentials-credentials.yaml
```

You can verify that the IAM service accounts are created by querying GCP. For more information, refer to [GCP documentation on listing IAM service accounts](#).

9.6.7.2.3. Incorporating the Cloud Credential Operator utility manifests

To implement short-term security credentials managed outside the cluster for individual components, you must move the manifest files that the Cloud Credential Operator utility (**ccoctl**) created to the correct directories for the installation program.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have configured the Cloud Credential Operator utility (**ccoctl**).
- You have created the cloud provider resources that are required for your cluster with the **ccoctl** utility.

Procedure

1. Add the following granular permissions to the GCP account that the installation program uses:

Example 9.37. Required GCP permissions

- `compute.machineTypes.list`
- `compute.regions.list`
- `compute.zones.list`
- `dns.changes.create`
- `dns.changes.get`
- `dns.managedZones.create`
- `dns.managedZones.delete`

- `dns.managedZones.get`
- `dns.managedZones.list`
- `dns.networks.bindPrivateDNSZone`
- `dns.resourceRecordSets.create`
- `dns.resourceRecordSets.delete`
- `dns.resourceRecordSets.list`

2. If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```
apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...
```

3. If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

4. Copy the manifests that the **ccoctl** utility generated to the **manifests** directory that the installation program created by running the following command:


```
$ cp /<path_to_ccoctl_output_dir>/manifests/* ./manifests/
```

5. Copy the private key that the **ccoctl** utility generated in the **tls** directory to the installation directory by running the following command:

```
$ cp -a /<path_to_ccoctl_output_dir>/tls .
```

9.6.8. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.

- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.

Procedure

1. Remove any existing GCP credentials that do not use the service account key for the GCP account that you configured for your cluster and that are stored in the following locations:

- The **GOOGLE_CREDENTIALS**, **GOOGLE_CLOUD_KEYFILE_JSON**, or **GKLOUD_KEYFILE_JSON** environment variables
- The `~/.gcp/osServiceAccount.json` file
- The **gcloud cli** default credentials

2. Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

- 1 For `<installation_directory>`, specify the location of your customized `./install-config.yaml` file.
- 2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

3. Optional: You can reduce the number of permissions for the service account that you used to install the cluster.

- If you assigned the **Owner** role to your service account, you can remove that role and replace it with the **Viewer** role.
- If you included the **Service Account Key Admin** role, you can remove it.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to `<installation_directory>/openshift_install.log`.



IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.


Example output

```
...
```

```

INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s

```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

9.6.9. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

-

9.6.10. Disabling the default OperatorHub catalog sources

Operator catalogs that source content provided by Red Hat and community projects are configured for OperatorHub by default during an OpenShift Container Platform installation. In a restricted network environment, you must disable the default catalogs as a cluster administrator.

Procedure

- Disable the sources for the default catalogs by adding **disableAllDefaultSources: true** to the **OperatorHub** object:

```
$ oc patch OperatorHub cluster --type json \
  -p '[{"op": "add", "path": "/spec/disableAllDefaultSources", "value": true}]'
```

TIP

Alternatively, you can use the web console to manage catalog sources. From the **Administration** → **Cluster Settings** → **Configuration** → **OperatorHub** page, click the **Sources** tab, where you can create, update, delete, disable, and enable individual sources.

9.6.11. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service

9.6.12. Next steps

- [Validate an installation](#).
- [Customize your cluster](#).
- [Configure image streams](#) for the Cluster Samples Operator and the **must-gather** tool.
- Learn how to [use Operator Lifecycle Manager \(OLM\) on restricted networks](#) .
- If the mirror registry that you used to install your cluster has a trusted CA, add it to the cluster by [configuring additional trust stores](#).
- If necessary, you can [opt out of remote health reporting](#) .
- If necessary, see [Registering your disconnected cluster](#)

9.7. INSTALLING A CLUSTER ON GCP INTO AN EXISTING VPC

In OpenShift Container Platform version 4.16, you can install a cluster into an existing Virtual Private Cloud (VPC) on Google Cloud Platform (GCP). The installation program provisions the rest of the required infrastructure, which you can further customize. To customize the installation, you modify parameters in the **install-config.yaml** file before you install the cluster.

9.7.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [configured a GCP project](#) to host the cluster.
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.

9.7.2. About using a custom VPC

In OpenShift Container Platform 4.16, you can deploy a cluster into existing subnets in an existing Virtual Private Cloud (VPC) in Google Cloud Platform (GCP). By deploying OpenShift Container Platform into an existing GCP VPC, you might be able to avoid limit constraints in new accounts or more easily abide by the operational constraints that your company's guidelines set. If you cannot obtain the infrastructure creation permissions that are required to create the VPC yourself, use this installation option. You must configure networking for the subnets.

9.7.2.1. Requirements for using your VPC

The union of the VPC CIDR block and the machine network CIDR must be non-empty. The subnets must be within the machine network.

The installation program does not create the following components:

- NAT gateways
- Subnets
- Route tables
- VPC network



NOTE

The installation program requires that you use the cloud-provided DNS server. Using a custom DNS server is not supported and causes the installation to fail.

9.7.2.2. VPC validation

To ensure that the subnets that you provide are suitable, the installation program confirms the following data:

- All the subnets that you specify exist.
- You provide one subnet for control-plane machines and one subnet for compute machines.

- The subnet's CIDRs belong to the machine CIDR that you specified.

9.7.2.3. Division of permissions

Some individuals can create different resource in your clouds than others. For example, you might be able to create application-specific items, like instances, buckets, and load balancers, but not networking-related components such as VPCs, subnets, or ingress rules.

9.7.2.4. Isolation between clusters

If you deploy OpenShift Container Platform to an existing network, the isolation of cluster services is reduced in the following ways:


- You can install multiple OpenShift Container Platform clusters in the same VPC.
- ICMP ingress is allowed to the entire network.
- TCP 22 ingress (SSH) is allowed to the entire network.
- Control plane TCP 6443 ingress (Kubernetes API) is allowed to the entire network.
- Control plane TCP 22623 ingress (MCS) is allowed to the entire network.

9.7.3. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.



IMPORTANT

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

9.7.4. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the `core` user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The **./openshift-install gather** command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.



NOTE

You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.



NOTE

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the **./openshift-install gather** command.

**NOTE**

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```

**NOTE**

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

9.7.5. Obtaining the installation program

Before you install OpenShift Container Platform, download the installation file on the host you are using for installation.

Prerequisites

- You have a computer that runs Linux or macOS, with 500 MB of local disk space.

Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select your infrastructure provider.

3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.



IMPORTANT

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.



IMPORTANT

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.

4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#). This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

9.7.6. Creating the installation configuration file

You can customize the OpenShift Container Platform cluster you install on Google Cloud Platform (GCP).

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create the **install-config.yaml** file.
 - a. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install create install-config --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the directory name to store the files that the installation program creates.


When specifying the directory:

- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
 - Use an empty directory. Some installation assets, such as bootstrap X.509 certificates, have short expiration intervals, therefore you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.
- b. At the prompts, provide the configuration details for your cloud:
- i. Optional: Select an SSH key to use to access your cluster machines.

**NOTE**

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- ii. Select **gcp** as the platform to target.
 - iii. If you have not configured the service account key for your GCP account on your computer, you must obtain it from GCP and paste the contents of the file or enter the absolute path to the file.
 - iv. Select the project ID to provision the cluster in. The default value is specified by the service account that you configured.
 - v. Select the region to deploy the cluster to.
 - vi. Select the base domain to deploy the cluster to. The base domain corresponds to the public DNS zone that you created for your cluster.
 - vii. Enter a descriptive name for your cluster.
2. Modify the **install-config.yaml** file. You can find more information about the available parameters in the "Installation configuration parameters" section.
 3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.

**IMPORTANT**

The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

Additional resources

- [Installation configuration parameters for GCP](#)

9.7.6.1. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 9.19. Minimum resource requirements

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Bootstrap	RHCOS	4	16 GB	100 GB	300
Control plane	RHCOS	4	16 GB	100 GB	300
Compute	RHCOS, RHEL 8.6 and later [3]	2	8 GB	100 GB	300

1. One vCPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio: (threads per core × cores) × sockets = vCPUs.
2. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes which require a 10 ms p99 fsync duration. Note that on many cloud platforms, storage size and IOPS scale together, so you might need to over-allocate storage volume to obtain sufficient performance.
3. As with all user-provisioned installations, if you choose to use RHEL compute machines in your cluster, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Use of RHEL 7 compute machines is deprecated and has been removed in OpenShift Container Platform 4.10 and later.

NOTE

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

9.7.6.2. Tested instance types for GCP

The following Google Cloud Platform instance types have been tested with OpenShift Container Platform.

Example 9.38. Machine series

- **A2**
- **A3**
- **C2**
- **C2D**
- **C3**
- **C3D**
- **E2**
- **M1**
- **N1**
- **N2**
- **N2D**
- **Tau T2D**

9.7.6.3. Tested instance types for GCP on 64-bit ARM infrastructures

The following Google Cloud Platform (GCP) 64-bit ARM instance types have been tested with OpenShift Container Platform.

Example 9.39. Machine series for 64-bit ARM machines

- **Tau T2A**

9.7.6.4. Using custom machine types

Using a custom machine type to install a OpenShift Container Platform cluster is supported.

Consider the following when using a custom machine type:

- Similar to predefined instance types, custom machine types must meet the minimum resource requirements for control plane and compute machines. For more information, see "Minimum resource requirements for cluster installation".
- The name of the custom machine type must adhere to the following syntax:
custom-`<number_of_cpus>-<amount_of_memory_in_mb>`

For example, **custom-6-20480**.

As part of the installation process, you specify the custom machine type in the **install-config.yaml** file.

Sample `install-config.yaml` file with a custom machine type

```
compute:
- architecture: amd64
  hyperthreading: Enabled
  name: worker
  platform:
    gcp:
      type: custom-6-20480
  replicas: 2
controlPlane:
  architecture: amd64
  hyperthreading: Enabled
  name: master
  platform:
    gcp:
      type: custom-6-20480
  replicas: 3
```

9.7.6.5. Enabling Shielded VMs

You can use Shielded VMs when installing your cluster. Shielded VMs have extra security features including secure boot, firmware and integrity monitoring, and rootkit detection. For more information, see Google's documentation on [Shielded VMs](#).



NOTE

Shielded VMs are currently not supported on clusters with 64-bit ARM infrastructures.

Prerequisites

- You have created an `install-config.yaml` file.

Procedure

- Use a text editor to edit the `install-config.yaml` file prior to deploying your cluster and add one of the following stanzas:

- To use shielded VMs for only control plane machines:

```
controlPlane:
  platform:
    gcp:
      secureBoot: Enabled
```

- To use shielded VMs for only compute machines:

```
compute:
- platform:
    gcp:
      secureBoot: Enabled
```

- To use shielded VMs for all machines:

■

```
platform:
  gcp:
    defaultMachinePlatform:
      secureBoot: Enabled
```

9.7.6.6. Enabling Confidential VMs

You can use Confidential VMs when installing your cluster. Confidential VMs encrypt data while it is being processed. For more information, see Google's documentation on [Confidential Computing](#). You can enable Confidential VMs and Shielded VMs at the same time, although they are not dependent on each other.



NOTE

Confidential VMs are currently not supported on 64-bit ARM architectures.

Prerequisites

- You have created an **install-config.yaml** file.

Procedure

- Use a text editor to edit the **install-config.yaml** file prior to deploying your cluster and add one of the following stanzas:
 - To use confidential VMs for only control plane machines:

```
controlPlane:
  platform:
    gcp:
      confidentialCompute: Enabled 1
      type: n2d-standard-8 2
      onHostMaintenance: Terminate 3
```

- Enable confidential VMs.
- Specify a machine type that supports Confidential VMs. Confidential VMs require the N2D or C2D series of machine types. For more information on supported machine types, see [Supported operating systems and machine types](#).
- Specify the behavior of the VM during a host maintenance event, such as a hardware or software update. For a machine that uses Confidential VM, this value must be set to **Terminate**, which stops the VM. Confidential VMs do not support live VM migration.

- To use confidential VMs for only compute machines:

```
compute:
  - platform:
      gcp:
        confidentialCompute: Enabled
        type: n2d-standard-8
        onHostMaintenance: Terminate
```

- c. To use confidential VMs for all machines:

```
platform:
  gcp:
    defaultMachinePlatform:
      confidentialCompute: Enabled
      type: n2d-standard-8
      onHostMaintenance: Terminate
```

9.7.6.7. Sample customized install-config.yaml file for GCP

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.



IMPORTANT

This sample YAML file is provided for reference only. You must obtain your **install-config.yaml** file by using the installation program and modify it.

```
apiVersion: v1
baseDomain: example.com 1
credentialsMode: Mint 2
controlPlane: 3 4
  hyperthreading: Enabled 5
  name: master
  platform:
    gcp:
      type: n2-standard-4
      zones:
      - us-central1-a
      - us-central1-c
    osDisk:
      diskType: pd-ssd
      diskSizeGB: 1024
      encryptionKey: 6
      kmsKey:
        name: worker-key
        keyRing: test-machine-keys
        location: global
        projectID: project-id
    tags: 7
      - control-plane-tag1
      - control-plane-tag2
    osImage: 8
      project: example-project-name
      name: example-image-name
  replicas: 3
compute: 9 10
- hyperthreading: Enabled 11
  name: worker
  platform:
    gcp:
      type: n2-standard-4
```

```

zones:
- us-central1-a
- us-central1-c
osDisk:
  diskType: pd-standard
  diskSizeGB: 128
  encryptionKey: 12
  kmsKey:
    name: worker-key
    keyRing: test-machine-keys
    location: global
    projectID: project-id
  tags: 13
  - compute-tag1
  - compute-tag2
  osImage: 14
    project: example-project-name
    name: example-image-name
replicas: 3
metadata:
  name: test-cluster 15
networking:
  clusterNetwork:
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  machineNetwork:
  - cidr: 10.0.0.0/16
  networkType: OVNKubernetes 16
  serviceNetwork:
  - 172.30.0.0/16
platform:
  gcp:
  projectID: openshift-production 17
  region: us-central1 18
  defaultMachinePlatform:
  tags: 19
  - global-tag1
  - global-tag2
  osImage: 20
  project: example-project-name
  name: example-image-name
  network: existing_vpc 21
  controlPlaneSubnet: control_plane_subnet 22
  computeSubnet: compute_subnet 23
pullSecret: '{"auths": ...}' 24
fips: false 25
sshKey: ssh-ed25519 AAAA... 26

```

1 15 17 18 24 Required. The installation program prompts you for this value.

2 Optional: Add this parameter to force the Cloud Credential Operator (CCO) to use the specified mode. By default, the CCO uses the root credentials in the **kube-system** namespace to dynamically try to determine the capabilities of the credentials. For details about CCO modes, see the "About the Cloud Credential Operator" section in the *Authentication and authorization* guide.

- 3 9 If you do not provide these parameters and values, the installation program provides the default value.
- 4 10 The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.
- 5 11 Whether to enable or disable simultaneous multithreading, or **hyperthreading**. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores. You can disable it by setting the parameter value to **Disabled**. If you disable simultaneous multithreading in some cluster machines, you must disable it in all cluster machines.



IMPORTANT

If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance. Use larger machine types, such as **n1-standard-8**, for your machines if you disable simultaneous multithreading.

- 6 12 Optional: The custom encryption key section to encrypt both virtual machines and persistent volumes. Your default compute service account must have the permissions granted to use your KMS key and have the correct IAM role assigned. The default service account name follows the **service-<project_number>@compute-system.iam.gserviceaccount.com** pattern. For more information about granting the correct permissions for your service account, see "Machine management" → "Creating compute machine sets" → "Creating a compute machine set on GCP".
- 7 13 19 Optional: A set of network tags to apply to the control plane or compute machine sets. The **platform.gcp.defaultMachinePlatform.tags** parameter will apply to both control plane and compute machines. If the **compute.platform.gcp.tags** or **controlPlane.platform.gcp.tags** parameters are set, they override the **platform.gcp.defaultMachinePlatform.tags** parameter.
- 8 14 20 Optional: A custom Red Hat Enterprise Linux CoreOS (RHCOS) that should be used to boot control plane and compute machines. The **project** and **name** parameters under **platform.gcp.defaultMachinePlatform.osImage** apply to both control plane and compute machines. If the **project** and **name** parameters under **controlPlane.platform.gcp.osImage** or **compute.platform.gcp.osImage** are set, they override the **platform.gcp.defaultMachinePlatform.osImage** parameters.
- 16 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 21 Specify the name of an existing VPC.
- 22 Specify the name of the existing subnet to deploy the control plane machines to. The subnet must belong to the VPC that you specified.
- 23 Specify the name of the existing subnet to deploy the compute machines to. The subnet must belong to the VPC that you specified.
- 25 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.



IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 26 You can optionally provide the **sshKey** value that you use to access the machines in your cluster.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

Additional resources

- [Enabling customer-managed encryption keys for a compute machine set](#)

9.7.6.8. Create an Ingress Controller with global access on GCP

You can create an Ingress Controller that has global access to a Google Cloud Platform (GCP) cluster. Global access is only available to Ingress Controllers using internal load balancers.

Prerequisites

- You created the **install-config.yaml** and complete any modifications to it.

Procedure

Create an Ingress Controller with global access on a new GCP cluster.

1. Change to the directory that contains the installation program and create a manifest file:

```
$ ./openshift-install create manifests --dir <installation_directory> 1
```

- 1 For **<installation_directory>**, specify the name of the directory that contains the **install-config.yaml** file for your cluster.

2. Create a file that is named **cluster-ingress-default-ingresscontroller.yaml** in the **<installation_directory>/manifests/** directory:

```
$ touch <installation_directory>/manifests/cluster-ingress-default-ingresscontroller.yaml 1
```

- 1 For **<installation_directory>**, specify the directory name that contains the **manifests/** directory for your cluster.

After creating the file, several network configuration files are in the **manifests/** directory, as shown:

```
$ ls <installation_directory>/manifests/cluster-ingress-default-ingresscontroller.yaml
```

Example output

```
cluster-ingress-default-ingresscontroller.yaml
```

3. Open the **cluster-ingress-default-ingresscontroller.yaml** file in an editor and enter a custom resource (CR) that describes the Operator configuration you want:

Sample `clientAccess` configuration to Global

```
apiVersion: operator.openshift.io/v1
kind: IngressController
metadata:
  name: default
  namespace: openshift-ingress-operator
spec:
  endpointPublishingStrategy:
    loadBalancer:
      providerParameters:
        gcp:
          clientAccess: Global 1
          type: GCP
          scope: Internal 2
          type: LoadBalancerService
```

- 1** Set `gcp.clientAccess` to **Global**.
- 2** Global access is only available to Ingress Controllers using internal load balancers.

9.7.6.9. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.



NOTE

The installation program does not support the proxy **readinessEndpoints** field.

**NOTE**

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.

**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

9.7.7. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.

**IMPORTANT**

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**. To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.



NOTE

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your PATH.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

9.7.8. Alternatives to storing administrator-level secrets in the kube-system project

By default, administrator secrets are stored in the **kube-system** project. If you configured the **credentialsMode** parameter in the **install-config.yaml** file to **Manual**, you must use one of the following alternatives:

- To manage long-term cloud credentials manually, follow the procedure in [Manually creating long-term credentials](#).
- To implement short-term credentials that are managed outside the cluster for individual components, follow the procedures in [Configuring a GCP cluster to use short-term credentials](#).

9.7.8.1. Manually creating long-term credentials

The Cloud Credential Operator (CCO) can be put into manual mode prior to installation in environments where the cloud identity and access management (IAM) APIs are not reachable, or the administrator prefers not to store an administrator-level credential secret in the cluster **kube-system** namespace.

Procedure

1. Add the following granular permissions to the GCP account that the installation program uses:

Example 9.40. Required GCP permissions

- compute.machineTypes.list
- compute.regions.list
- compute.zones.list
- dns.changes.create
- dns.changes.get
- dns.managedZones.create
- dns.managedZones.delete
- dns.managedZones.get
- dns.managedZones.list
- dns.networks.bindPrivateDNSZone
- dns.resourceRecordSets.create
- dns.resourceRecordSets.delete
- dns.resourceRecordSets.list

- If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```
apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...
```

- If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

- Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

- Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included 1 \
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml 2 \
  --to=<path_to_directory_for_credentials_requests> 3
```

- The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- Specify the location of the **install-config.yaml** file.
- Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

This command creates a YAML file for each **CredentialsRequest** object.

Sample CredentialsRequest object

```
apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
```

```

kind: GCPProviderSpec
predefinedRoles:
- roles/storage.admin
- roles/iam.serviceAccountUser
skipServiceCheck: true
...

```

6. Create YAML files for secrets in the **openshift-install** manifests directory that you generated previously. The secrets must be stored using the namespace and secret name defined in the **spec.secretRef** for each **CredentialsRequest** object.

Sample CredentialsRequest object with secrets

```

apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    ...
  secretRef:
    name: <component_secret>
    namespace: <component_namespace>
  ...

```

Sample Secret object

```

apiVersion: v1
kind: Secret
metadata:
  name: <component_secret>
  namespace: <component_namespace>
data:
  service_account.json: <base64_encoded_gcp_service_account_file>

```



IMPORTANT


Before upgrading a cluster that uses manually maintained credentials, you must ensure that the CCO is in an upgradeable state.

9.7.8.2. Configuring a GCP cluster to use short-term credentials

To install a cluster that is configured to use GCP Workload Identity, you must configure the CCO utility and create the required GCP resources for your cluster.

9.7.8.2.1. Configuring the Cloud Credential Operator utility

To create and manage cloud credentials from outside of the cluster when the Cloud Credential Operator (CCO) is operating in manual mode, extract and prepare the CCO utility (**ccoctl**) binary.



NOTE

The **ccoctl** utility is a Linux binary that must run in a Linux environment.

Prerequisites

- You have access to an OpenShift Container Platform account with cluster administrator access.
- You have installed the OpenShift CLI (**oc**).
- You have added one of the following authentication options to the GCP account that the installation program uses:
 - The **IAM Workload Identity Pool Admin** role.
 - The following granular permissions:

Example 9.41. Required GCP permissions

- `compute.projects.get`
- `iam.googleapis.com/workloadIdentityPoolProviders.create`
- `iam.googleapis.com/workloadIdentityPoolProviders.get`
- `iam.googleapis.com/workloadIdentityPools.create`
- `iam.googleapis.com/workloadIdentityPools.delete`
- `iam.googleapis.com/workloadIdentityPools.get`
- `iam.googleapis.com/workloadIdentityPools.undelete`
- `iam.roles.create`
- `iam.roles.delete`
- `iam.roles.list`
- `iam.roles.undelete`
- `iam.roles.update`
- `iam.serviceAccounts.create`
- `iam.serviceAccounts.delete`
- `iam.serviceAccounts.getIamPolicy`
- `iam.serviceAccounts.list`
- `iam.serviceAccounts.setIamPolicy`
- `iam.workloadIdentityPoolProviders.get`
- `iam.workloadIdentityPools.delete`
- `resourceManager.projects.get`

- resourcemanager.projects.getlamPolicy
- resourcemanager.projects.setlamPolicy
- storage.buckets.create
- storage.buckets.delete
- storage.buckets.get
- storage.buckets.getlamPolicy
- storage.buckets.setlamPolicy
- storage.objects.create
- storage.objects.delete
- storage.objects.list

Procedure

1. Set a variable for the OpenShift Container Platform release image by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

2. Obtain the CCO container image from the OpenShift Container Platform release image by running the following command:

```
$ CCO_IMAGE=$(oc adm release info --image-for='cloud-credential-operator'
$RELEASE_IMAGE -a ~/.pull-secret)
```



NOTE

Ensure that the architecture of the **\$RELEASE_IMAGE** matches the architecture of the environment in which you will use the **ccoctl** tool.

3. Extract the **ccoctl** binary from the CCO container image within the OpenShift Container Platform release image by running the following command:

```
$ oc image extract $CCO_IMAGE \
--file="/usr/bin/ccoctl.<rhel_version>" \
-a ~/.pull-secret
```

- 1 For **<rhel_version>**, specify the value that corresponds to the version of Red Hat Enterprise Linux (RHEL) that the host uses. If no value is specified, **ccoctl.rhel8** is used by default. The following values are valid:

- **rhel8**: Specify this value for hosts that use RHEL 8.
- **rhel9**: Specify this value for hosts that use RHEL 9.

- Change the permissions to make **ccoctl** executable by running the following command:

```
$ chmod 775 ccoctl.<rhel_version>
```

Verification

- To verify that **ccoctl** is ready to use, display the help file by running the following command:

```
$ ccoctl --help
```

Output of **ccoctl --help**

```
OpenShift credentials provisioning tool
```

```
Usage:
```

```
ccoctl [command]
```

```
Available Commands:
```

```
aws      Manage credentials objects for AWS cloud
azure    Manage credentials objects for Azure
gcp      Manage credentials objects for Google cloud
help     Help about any command
ibmcloud Manage credentials objects for IBM Cloud
nutanix  Manage credentials objects for Nutanix
```

```
Flags:
```

```
-h, --help  help for ccoctl
```

```
Use "ccoctl [command] --help" for more information about a command.
```

9.7.8.2.2. Creating GCP resources with the Cloud Credential Operator utility

You can use the **ccoctl gcp create-all** command to automate the creation of GCP resources.



NOTE

By default, **ccoctl** creates objects in the directory in which the commands are run. To create the objects in a different directory, use the **--output-dir** flag. This procedure uses **<path_to_ccoctl_output_dir>** to refer to this directory.

Prerequisites

You must have:

- Extracted and prepared the **ccoctl** binary.

Procedure

- Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

2. Extract the list of **CredentialsRequest** objects from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included \ 1
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \ 2
  --to=<path_to_directory_for_credentials_requests> \ 3
```

- 1 The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2 Specify the location of the **install-config.yaml** file.
- 3 Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.



NOTE

This command might take a few moments to run.

3. Use the **ccoctl** tool to process all **CredentialsRequest** objects by running the following command:

```
$ ccoctl gcp create-all \
  --name=<name> \ 1
  --region=<gcp_region> \ 2
  --project=<gcp_project_id> \ 3
  --credentials-requests-dir=<path_to_credentials_requests_directory> \ 4
```

- 1 Specify the user-defined name for all created GCP resources used for tracking.
- 2 Specify the GCP region in which cloud resources will be created.
- 3 Specify the GCP project ID in which cloud resources will be created.
- 4 Specify the directory containing the files of **CredentialsRequest** manifests to create GCP service accounts.



NOTE

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

Verification

- To verify that the OpenShift Container Platform secrets are created, list the files in the **<path_to_ccoctl_output_dir>/manifests** directory:

```
$ ls <path_to_ccoctl_output_dir>/manifests
```

Example output

```
cluster-authentication-02-config.yaml
openshift-cloud-controller-manager-gcp-ccm-cloud-credentials-credentials.yaml
openshift-cloud-credential-operator-cloud-credential-operator-gcp-ro-creds-credentials.yaml
openshift-cloud-network-config-controller-cloud-credentials-credentials.yaml
openshift-cluster-api-capg-manager-bootstrap-credentials-credentials.yaml
openshift-cluster-csi-drivers-gcp-pd-cloud-credentials-credentials.yaml
openshift-image-registry-installer-cloud-credentials-credentials.yaml
openshift-ingress-operator-cloud-credentials-credentials.yaml
openshift-machine-api-gcp-cloud-credentials-credentials.yaml
```

You can verify that the IAM service accounts are created by querying GCP. For more information, refer to GCP documentation on listing IAM service accounts.

9.7.8.2.3. Incorporating the Cloud Credential Operator utility manifests

To implement short-term security credentials managed outside the cluster for individual components, you must move the manifest files that the Cloud Credential Operator utility (**ccoctl**) created to the correct directories for the installation program.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have configured the Cloud Credential Operator utility (**ccoctl**).
- You have created the cloud provider resources that are required for your cluster with the **ccoctl** utility.

Procedure

1. Add the following granular permissions to the GCP account that the installation program uses:

Example 9.42. Required GCP permissions

- compute.machineTypes.list
- compute.regions.list
- compute.zones.list
- dns.changes.create
- dns.changes.get
- dns.managedZones.create
- dns.managedZones.delete
- dns.managedZones.get
- dns.managedZones.list
- dns.networks.bindPrivateDNSZone

- `dns.resourceRecordSets.create`
- `dns.resourceRecordSets.delete`
- `dns.resourceRecordSets.list`

2. If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```
apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...
```

3. If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

4. Copy the manifests that the **ccoctl** utility generated to the **manifests** directory that the installation program created by running the following command:

```
$ cp /<path_to_ccoctl_output_dir>/manifests/* ./manifests/
```

5. Copy the private key that the **ccoctl** utility generated in the **tls** directory to the installation directory by running the following command:

```
$ cp -a /<path_to_ccoctl_output_dir>/tls .
```

9.7.9. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.

Procedure

1. Remove any existing GCP credentials that do not use the service account key for the GCP account that you configured for your cluster and that are stored in the following locations:
 - The **GOOGLE_CREDENTIALS**, **GOOGLE_CLOUD_KEYFILE_JSON**, or **GKLOUD_KEYFILE_JSON** environment variables
 - The `~/.gcp/osServiceAccount.json` file
 - The **gcloud cli** default credentials
2. Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

1 For `<installation_directory>`, specify the location of your customized `./install-config.yaml` file.


2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

3. Optional: You can reduce the number of permissions for the service account that you used to install the cluster.
 - If you assigned the **Owner** role to your service account, you can remove that role and replace it with the **Viewer** role.
 - If you included the **Service Account Key Admin** role, you can remove it.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to `<installation_directory>/openshift_install.log`.



IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

9.7.10. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

Additional resources

- See [Accessing the web console](#) for more details about accessing and understanding the OpenShift Container Platform web console.

9.7.11. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service

9.7.12. Next steps

- [Customize your cluster](#).
- If necessary, you can [opt out of remote health reporting](#).

9.8. INSTALLING A CLUSTER ON GCP INTO A SHARED VPC

In OpenShift Container Platform version 4.16, you can install a cluster into a shared Virtual Private Cloud (VPC) on Google Cloud Platform (GCP). In this installation method, the cluster is configured to use a VPC from a different GCP project. A shared VPC enables an organization to connect resources from multiple projects to a common VPC network. You can communicate within the organization securely and efficiently by using internal IP addresses from that network. For more information about shared VPC, see [Shared VPC overview in the GCP documentation](#).

The installation program provisions the rest of the required infrastructure, which you can further customize. To customize the installation, you modify parameters in the `install-config.yaml` file before you install the cluster.

9.8.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.
- You have a GCP host project which contains a shared VPC network.
- You [configured a GCP project](#) to host the cluster. This project, known as the service project, must be attached to the host project. For more information, see [Attaching service projects in the GCP documentation](#).
- You have a GCP service account that has the [required GCP permissions](#) in both the host and service projects.

9.8.2. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.



IMPORTANT

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

9.8.3. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.



NOTE

You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.

**NOTE**

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.

**NOTE**

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```

**NOTE**

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

9.8.4. Obtaining the installation program

Before you install OpenShift Container Platform, download the installation file on the host you are using for installation.

Prerequisites

- You have a computer that runs Linux or macOS, with 500 MB of local disk space.


Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select your infrastructure provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.



IMPORTANT

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.



IMPORTANT

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.

4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#) . This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

9.8.5. Creating the installation files for GCP

To install OpenShift Container Platform on Google Cloud Platform (GCP) into a shared VPC, you must generate the **install-config.yaml** file and modify it so that the cluster uses the correct VPC networks, DNS zones, and project names.

9.8.5.1. Manually creating the installation configuration file

Installing the cluster requires that you manually create the installation configuration file.


Prerequisites

- You have an SSH public key on your local machine to provide to the installation program. The key will be used for SSH authentication onto your cluster nodes for debugging and disaster recovery.
- You have obtained the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create an installation directory to store your required installation assets in:

```
$ mkdir <installation_directory>
```



IMPORTANT

You must create a directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

2. Customize the sample **install-config.yaml** file template that is provided and save it in the **<installation_directory>**.



NOTE

You must name this configuration file **install-config.yaml**.

3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the next step of the installation process. You must back it up now.

Additional resources

- [Installation configuration parameters for GCP](#)

9.8.5.2. Enabling Shielded VMs

You can use Shielded VMs when installing your cluster. Shielded VMs have extra security features including secure boot, firmware and integrity monitoring, and rootkit detection. For more information, see Google's documentation on [Shielded VMs](#).

**NOTE**

Shielded VMs are currently not supported on clusters with 64-bit ARM infrastructures.

Prerequisites

- You have created an **install-config.yaml** file.

Procedure

- Use a text editor to edit the **install-config.yaml** file prior to deploying your cluster and add one of the following stanzas:
 - To use shielded VMs for only control plane machines:

```
controlPlane:
  platform:
    gcp:
      secureBoot: Enabled
```

- To use shielded VMs for only compute machines:

```
compute:
- platform:
  gcp:
    secureBoot: Enabled
```

- To use shielded VMs for all machines:

```
platform:
  gcp:
    defaultMachinePlatform:
      secureBoot: Enabled
```

9.8.5.3. Enabling Confidential VMs

You can use Confidential VMs when installing your cluster. Confidential VMs encrypt data while it is being processed. For more information, see Google's documentation on [Confidential Computing](#). You can enable Confidential VMs and Shielded VMs at the same time, although they are not dependent on each other.

**NOTE**

Confidential VMs are currently not supported on 64-bit ARM architectures.

Prerequisites

- You have created an **install-config.yaml** file.

Procedure

- Use a text editor to edit the **install-config.yaml** file prior to deploying your cluster and add one of the following stanzas:

- a. To use confidential VMs for only control plane machines:

```
controlPlane:
  platform:
    gcp:
      confidentialCompute: Enabled 1
      type: n2d-standard-8 2
      onHostMaintenance: Terminate 3
```

- 1 Enable confidential VMs.
- 2 Specify a machine type that supports Confidential VMs. Confidential VMs require the N2D or C2D series of machine types. For more information on supported machine types, see [Supported operating systems and machine types](#).
- 3 Specify the behavior of the VM during a host maintenance event, such as a hardware or software update. For a machine that uses Confidential VM, this value must be set to **Terminate**, which stops the VM. Confidential VMs do not support live VM migration.

- b. To use confidential VMs for only compute machines:

```
compute:
  - platform:
      gcp:
        confidentialCompute: Enabled
        type: n2d-standard-8
        onHostMaintenance: Terminate
```

- c. To use confidential VMs for all machines:

```
platform:
  gcp:
    defaultMachinePlatform:
      confidentialCompute: Enabled
      type: n2d-standard-8
      onHostMaintenance: Terminate
```

9.8.5.4. Sample customized install-config.yaml file for shared VPC installation

There are several configuration parameters which are required to install OpenShift Container Platform on GCP using a shared VPC. The following is a sample **install-config.yaml** file which demonstrates these fields.



IMPORTANT

This sample YAML file is provided for reference only. You must modify this file with the correct values for your environment and cluster.

```
apiVersion: v1
baseDomain: example.com
credentialsMode: Passthrough 1
metadata:
```

```

name: cluster_name
platform:
  gcp:
    computeSubnet: shared-vpc-subnet-1 2
    controlPlaneSubnet: shared-vpc-subnet-2 3
    network: shared-vpc 4
    networkProjectID: host-project-name 5
    projectID: service-project-name 6
    region: us-east1
    defaultMachinePlatform:
      tags: 7
      - global-tag1
  controlPlane:
    name: master
    platform:
      gcp:
        tags: 8
        - control-plane-tag1
        type: n2-standard-4
        zones:
          - us-central1-a
          - us-central1-c
    replicas: 3
  compute:
    - name: worker
    platform:
      gcp:
        tags: 9
        - compute-tag1
        type: n2-standard-4
        zones:
          - us-central1-a
          - us-central1-c
    replicas: 3
  networking:
    clusterNetwork:
      - cidr: 10.128.0.0/14
      hostPrefix: 23
    machineNetwork:
      - cidr: 10.0.0.0/16
  pullSecret: '{"auths": ...}'
  sshKey: ssh-ed25519 AAAA... 10

```

- 1 **credentialsMode** must be set to **Passthrough** or **Manual**. See the "Prerequisites" section for the required GCP permissions that your service account must have.
- 2 The name of the subnet in the shared VPC for compute machines to use.
- 3 The name of the subnet in the shared VPC for control plane machines to use.
- 4 The name of the shared VPC.
- 5 The name of the host project where the shared VPC exists.
- 6 The name of the GCP project where you want to install the cluster.

- 7 8 9 Optional. One or more network tags to apply to compute machines, control plane machines, or all machines.
- 10 You can optionally provide the **sshKey** value that you use to access the machines in your cluster.

9.8.5.5. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network

- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.

**NOTE**

The installation program does not support the proxy **readinessEndpoints** field.

**NOTE**

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.

**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

9.8.6. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.

**IMPORTANT**

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.

4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**.
To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.

**NOTE**

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your PATH.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

9.8.7. Alternatives to storing administrator-level secrets in the kube-system project

By default, administrator secrets are stored in the **kube-system** project. If you configured the **credentialsMode** parameter in the **install-config.yaml** file to **Manual**, you must use one of the following alternatives:

- To manage long-term cloud credentials manually, follow the procedure in [Manually creating long-term credentials](#).
- To implement short-term credentials that are managed outside the cluster for individual components, follow the procedures in [Configuring a GCP cluster to use short-term credentials](#).

9.8.7.1. Manually creating long-term credentials

The Cloud Credential Operator (CCO) can be put into manual mode prior to installation in environments where the cloud identity and access management (IAM) APIs are not reachable, or the administrator prefers not to store an administrator-level credential secret in the cluster **kube-system** namespace.

Procedure

1. Add the following granular permissions to the GCP account that the installation program uses:

Example 9.43. Required GCP permissions

- compute.machineTypes.list
- compute.regions.list
- compute.zones.list
- dns.changes.create
- dns.changes.get
- dns.managedZones.create
- dns.managedZones.delete

- `dns.managedZones.get`
- `dns.managedZones.list`
- `dns.networks.bindPrivateDNSZone`
- `dns.resourceRecordSets.create`
- `dns.resourceRecordSets.delete`
- `dns.resourceRecordSets.list`

2. If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```
apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...
```

3. If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

4. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

5. Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included 1 \
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml 2 \
  --to=<path_to_directory_for_credentials_requests> 3
```

- 1** The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2** Specify the location of the **install-config.yaml** file.
- 3** Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

This command creates a YAML file for each **CredentialsRequest** object.

Sample `CredentialsRequest` object

```

apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: GCProviderSpec
    predefinedRoles:
      - roles/storage.admin
      - roles/iam.serviceAccountUser
    skipServiceCheck: true
  ...

```

6. Create YAML files for secrets in the **openshift-install** manifests directory that you generated previously. The secrets must be stored using the namespace and secret name defined in the **spec.secretRef** for each **CredentialsRequest** object.

Sample `CredentialsRequest` object with secrets

```

apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    ...
  secretRef:
    name: <component_secret>
    namespace: <component_namespace>
  ...

```

Sample `Secret` object

```

apiVersion: v1
kind: Secret
metadata:
  name: <component_secret>
  namespace: <component_namespace>
data:
  service_account.json: <base64_encoded_gcp_service_account_file>

```



IMPORTANT

Before upgrading a cluster that uses manually maintained credentials, you must ensure that the CCO is in an upgradeable state.

9.8.7.2. Configuring a GCP cluster to use short-term credentials

To install a cluster that is configured to use GCP Workload Identity, you must configure the CCO utility and create the required GCP resources for your cluster.

9.8.7.2.1. Configuring the Cloud Credential Operator utility

To create and manage cloud credentials from outside of the cluster when the Cloud Credential Operator (CCO) is operating in manual mode, extract and prepare the CCO utility (**ccoctl**) binary.



NOTE

The **ccoctl** utility is a Linux binary that must run in a Linux environment.

Prerequisites

- You have access to an OpenShift Container Platform account with cluster administrator access.
- You have installed the OpenShift CLI (**oc**).
- You have added one of the following authentication options to the GCP account that the installation program uses:
 - The **IAM Workload Identity Pool Admin** role.
 - The following granular permissions:

Example 9.44. Required GCP permissions

- `compute.projects.get`
- `iam.googleapis.com/workloadIdentityPoolProviders.create`
- `iam.googleapis.com/workloadIdentityPoolProviders.get`
- `iam.googleapis.com/workloadIdentityPools.create`
- `iam.googleapis.com/workloadIdentityPools.delete`
- `iam.googleapis.com/workloadIdentityPools.get`
- `iam.googleapis.com/workloadIdentityPools.undelete`
- `iam.roles.create`
- `iam.roles.delete`
- `iam.roles.list`
- `iam.roles.undelete`
- `iam.roles.update`
- `iam.serviceAccounts.create`
- `iam.serviceAccounts.delete`
- `iam.serviceAccounts.getIamPolicy`

- iam.serviceAccounts.list
- iam.serviceAccounts.setIamPolicy
- iam.workloadIdentityPoolProviders.get
- iam.workloadIdentityPools.delete
- resourceManager.projects.get
- resourceManager.projects.getIamPolicy
- resourceManager.projects.setIamPolicy
- storage.buckets.create
- storage.buckets.delete
- storage.buckets.get
- storage.buckets.getIamPolicy
- storage.buckets.setIamPolicy
- storage.objects.create
- storage.objects.delete
- storage.objects.list


Procedure

1. Set a variable for the OpenShift Container Platform release image by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

2. Obtain the CCO container image from the OpenShift Container Platform release image by running the following command:

```
$ CCO_IMAGE=$(oc adm release info --image-for='cloud-credential-operator'
$RELEASE_IMAGE -a ~/.pull-secret)
```



NOTE

Ensure that the architecture of the **\$RELEASE_IMAGE** matches the architecture of the environment in which you will use the **ccoctl** tool.

3. Extract the **ccoctl** binary from the CCO container image within the OpenShift Container Platform release image by running the following command:

```
$ oc image extract $CCO_IMAGE \
--file="/usr/bin/ccoctl.<rhel_version>" \
-a ~/.pull-secret
```


■

- 1 For **<rhel_version>**, specify the value that corresponds to the version of Red Hat Enterprise Linux (RHEL) that the host uses. If no value is specified, **ccoctl.rhel8** is used by default. The following values are valid:

- **rhel8**: Specify this value for hosts that use RHEL 8.
- **rhel9**: Specify this value for hosts that use RHEL 9.

4. Change the permissions to make **ccoctl** executable by running the following command:

```
$ chmod 775 ccoctl.<rhel_version>
```

Verification

- To verify that **ccoctl** is ready to use, display the help file by running the following command:

```
$ ccoctl --help
```

Output of **ccoctl --help**

```
OpenShift credentials provisioning tool
```

```
Usage:
```

```
ccoctl [command]
```

```
Available Commands:
```

```
aws      Manage credentials objects for AWS cloud
azure    Manage credentials objects for Azure
gcp      Manage credentials objects for Google cloud
help     Help about any command
ibmcloud Manage credentials objects for IBM Cloud
nutanix  Manage credentials objects for Nutanix
```

```
Flags:
```

```
-h, --help  help for ccoctl
```

```
Use "ccoctl [command] --help" for more information about a command.
```

9.8.7.2.2. Creating GCP resources with the Cloud Credential Operator utility

You can use the **ccoctl gcp create-all** command to automate the creation of GCP resources.



NOTE

By default, **ccoctl** creates objects in the directory in which the commands are run. To create the objects in a different directory, use the **--output-dir** flag. This procedure uses **<path_to_ccoctl_output_dir>** to refer to this directory.

Prerequisites

You must have:

- Extracted and prepared the **ccoctl** binary.

Procedure

1. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

2. Extract the list of **CredentialsRequest** objects from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included \1 \
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \2 \
  --to=<path_to_directory_for_credentials_requests> \3
```

- 1 The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2 Specify the location of the **install-config.yaml** file.
- 3 Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.




NOTE

This command might take a few moments to run.

3. Use the **ccoctl** tool to process all **CredentialsRequest** objects by running the following command:

```
$ ccoctl gcp create-all \
  --name=<name> \1 \
  --region=<gcp_region> \2 \
  --project=<gcp_project_id> \3 \
  --credentials-requests-dir=<path_to_credentials_requests_directory> \4
```

- 1 Specify the user-defined name for all created GCP resources used for tracking.
- 2 Specify the GCP region in which cloud resources will be created.
- 3 Specify the GCP project ID in which cloud resources will be created.
- 4 Specify the directory containing the files of **CredentialsRequest** manifests to create GCP service accounts.



NOTE

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

Verification

- To verify that the OpenShift Container Platform secrets are created, list the files in the `<path_to_ccoctl_output_dir>/manifests` directory:

```
$ ls <path_to_ccoctl_output_dir>/manifests
```

Example output

```
cluster-authentication-02-config.yaml
openshift-cloud-controller-manager-gcp-ccm-cloud-credentials-credentials.yaml
openshift-cloud-credential-operator-cloud-credential-operator-gcp-ro-creds-credentials.yaml
openshift-cloud-network-config-controller-cloud-credentials-credentials.yaml
openshift-cluster-api-capg-manager-bootstrap-credentials-credentials.yaml
openshift-cluster-csi-drivers-gcp-pd-cloud-credentials-credentials.yaml
openshift-image-registry-installer-cloud-credentials-credentials.yaml
openshift-ingress-operator-cloud-credentials-credentials.yaml
openshift-machine-api-gcp-cloud-credentials-credentials.yaml
```

You can verify that the IAM service accounts are created by querying GCP. For more information, refer to GCP documentation on listing IAM service accounts.

9.8.7.2.3. Incorporating the Cloud Credential Operator utility manifests

To implement short-term security credentials managed outside the cluster for individual components, you must move the manifest files that the Cloud Credential Operator utility (**ccoctl**) created to the correct directories for the installation program.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have configured the Cloud Credential Operator utility (**ccoctl**).
- You have created the cloud provider resources that are required for your cluster with the **ccoctl** utility.

Procedure

- Add the following granular permissions to the GCP account that the installation program uses:

Example 9.45. Required GCP permissions

- compute.machineTypes.list
- compute.regions.list
- compute.zones.list

- `dns.changes.create`
- `dns.changes.get`
- `dns.managedZones.create`
- `dns.managedZones.delete`
- `dns.managedZones.get`
- `dns.managedZones.list`
- `dns.networks.bindPrivateDNSZone`
- `dns.resourceRecordSets.create`
- `dns.resourceRecordSets.delete`
- `dns.resourceRecordSets.list`

2. If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```
apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...
```

3. If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

4. Copy the manifests that the **ccoctl** utility generated to the **manifests** directory that the installation program created by running the following command:


```
$ cp /<path_to_ccoctl_output_dir>/manifests/* ./manifests/
```

5. Copy the private key that the **ccoctl** utility generated in the **tls** directory to the installation directory by running the following command:

```
$ cp -a /<path_to_ccoctl_output_dir>/tls .
```

9.8.8. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.

Procedure

1. Remove any existing GCP credentials that do not use the service account key for the GCP account that you configured for your cluster and that are stored in the following locations:
 - The **GOOGLE_CREDENTIALS**, **GOOGLE_CLOUD_KEYFILE_JSON**, or **GCPLOUD_KEYFILE_JSON** environment variables
 - The `~/.gcp/osServiceAccount.json` file
 - The **gcloud cli** default credentials

2. Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```


- 1 For **<installation_directory>**, specify the location of your customized `./install-config.yaml` file.
- 2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

3. Optional: You can reduce the number of permissions for the service account that you used to install the cluster.
 - If you assigned the **Owner** role to your service account, you can remove that role and replace it with the **Viewer** role.
 - If you included the **Service Account Key Admin** role, you can remove it.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to `<installation_directory>/openshift_install.log`.



IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```

...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s

```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

9.8.9. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

Additional resources

- See [Accessing the web console](#) for more details about accessing and understanding the OpenShift Container Platform web console.

9.8.10. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service

9.8.11. Next steps

- [Customize your cluster](#).
- If necessary, you can [opt out of remote health reporting](#).

9.9. INSTALLING A PRIVATE CLUSTER ON GCP

In OpenShift Container Platform version 4.16, you can install a private cluster into an existing VPC on Google Cloud Platform (GCP). The installation program provisions the rest of the required infrastructure, which you can further customize. To customize the installation, you modify parameters in the **install-config.yaml** file before you install the cluster.

9.9.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [configured a GCP project](#) to host the cluster.
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.

9.9.2. Private clusters

You can deploy a private OpenShift Container Platform cluster that does not expose external endpoints. Private clusters are accessible from only an internal network and are not visible to the internet.

By default, OpenShift Container Platform is provisioned to use publicly-accessible DNS and endpoints. A private cluster sets the DNS, Ingress Controller, and API server to private when you deploy your cluster. This means that the cluster resources are only accessible from your internal network and are not visible to the internet.



IMPORTANT

If the cluster has any public subnets, load balancer services created by administrators might be publicly accessible. To ensure cluster security, verify that these services are explicitly annotated as private.

To deploy a private cluster, you must:

- Use existing networking that meets your requirements. Your cluster resources might be shared between other clusters on the network.
- Deploy from a machine that has access to:
 - The API services for the cloud to which you provision.
 - The hosts on the network that you provision.
 - The internet to obtain installation media.

You can use any machine that meets these access requirements and follows your company's guidelines. For example, this machine can be a bastion host on your cloud network or a machine that has access to the network through a VPN.

9.9.2.1. Private clusters in GCP

To create a private cluster on Google Cloud Platform (GCP), you must provide an existing private VPC and subnets to host the cluster. The installation program must also be able to resolve the DNS records that the cluster requires. The installation program configures the Ingress Operator and API server for only internal traffic.

The cluster still requires access to internet to access the GCP APIs.

The following items are not required or created when you install a private cluster:

- Public subnets
- Public network load balancers, which support public ingress
- A public DNS zone that matches the **baseDomain** for the cluster

The installation program does use the **baseDomain** that you specify to create a private DNS zone and the required records for the cluster. The cluster is configured so that the Operators do not create public records for the cluster and all cluster machines are placed in the private subnets that you specify.

Because it is not possible to limit access to external load balancers based on source tags, the private cluster uses only internal load balancers to allow access to internal instances.

The internal load balancer relies on instance groups rather than the target pools that the network load balancers use. The installation program creates instance groups for each zone, even if there is no instance in that group.

- The cluster IP address is internal only.
- One forwarding rule manages both the Kubernetes API and machine config server ports.
- The backend service is comprised of each zone's instance group and, while it exists, the bootstrap instance group.
- The firewall uses a single rule that is based on only internal source ranges.

9.9.2.1.1. Limitations

No health check for the Machine config server, `/healthz`, runs because of a difference in load balancer functionality. Two internal load balancers cannot share a single IP address, but two network load balancers can share a single external IP address. Instead, the health of an instance is determined entirely by the `/readyz` check on port 6443.

9.9.3. About using a custom VPC

In OpenShift Container Platform 4.16, you can deploy a cluster into an existing VPC in Google Cloud Platform (GCP). If you do, you must also use existing subnets within the VPC and routing rules.

By deploying OpenShift Container Platform into an existing GCP VPC, you might be able to avoid limit constraints in new accounts or more easily abide by the operational constraints that your company's guidelines set. This is a good option to use if you cannot obtain the infrastructure creation permissions that are required to create the VPC yourself.

9.9.3.1. Requirements for using your VPC

The installation program will no longer create the following components:

- VPC
- Subnets
- Cloud router
- Cloud NAT
- NAT IP addresses

If you use a custom VPC, you must correctly configure it and its subnets for the installation program and the cluster to use. The installation program cannot subdivide network ranges for the cluster to use, set route tables for the subnets, or set VPC options like DHCP, so you must do so before you install the cluster.

Your VPC and subnets must meet the following characteristics:

- The VPC must be in the same GCP project that you deploy the OpenShift Container Platform cluster to.
- To allow access to the internet from the control plane and compute machines, you must configure cloud NAT on the subnets to allow egress to it. These machines do not have a public

address. Even if you do not require access to the internet, you must allow egress to the VPC network to obtain the installation program and images. Because multiple cloud NATs cannot be configured on the shared subnets, the installation program cannot configure it.

To ensure that the subnets that you provide are suitable, the installation program confirms the following data:

- All the subnets that you specify exist and belong to the VPC that you specified.
- The subnet CIDRs belong to the machine CIDR.
- You must provide a subnet to deploy the cluster control plane and compute machines to. You can use the same subnet for both machine types.

If you destroy a cluster that uses an existing VPC, the VPC is not deleted.

9.9.3.2. Division of permissions

Starting with OpenShift Container Platform 4.3, you do not need all of the permissions that are required for an installation program-provisioned infrastructure cluster to deploy a cluster. This change mimics the division of permissions that you might have at your company: some individuals can create different resources in your clouds than others. For example, you might be able to create application-specific items, like instances, buckets, and load balancers, but not networking-related components such as VPCs, subnets, or Ingress rules.

The GCP credentials that you use when you create your cluster do not need the networking permissions that are required to make VPCs and core networking components within the VPC, such as subnets, routing tables, internet gateways, NAT, and VPN. You still need permission to make the application resources that the machines within the cluster require, such as load balancers, security groups, storage, and nodes.

9.9.3.3. Isolation between clusters

If you deploy OpenShift Container Platform to an existing network, the isolation of cluster services is preserved by firewall rules that reference the machines in your cluster by the cluster's infrastructure ID. Only traffic within the cluster is allowed.

If you deploy multiple clusters to the same VPC, the following components might share access between clusters:

- The API, which is globally available with an external publishing strategy or available throughout the network in an internal publishing strategy
- Debugging tools, such as ports on VM instances that are open to the machine CIDR for SSH and ICMP access

9.9.4. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.

- Access [Quay.io](https://quay.io) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.



IMPORTANT

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

9.9.5. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.



NOTE

You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.

**NOTE**

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.

**NOTE**

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```

**NOTE**

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

9.9.6. Obtaining the installation program

Before you install OpenShift Container Platform, download the installation file on the host you are using for installation.

Prerequisites

- You have a computer that runs Linux or macOS, with 500 MB of local disk space.


Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select your infrastructure provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.



IMPORTANT

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.



IMPORTANT

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.

4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#) . This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

9.9.7. Manually creating the installation configuration file

Installing the cluster requires that you manually create the installation configuration file.


Prerequisites

- You have an SSH public key on your local machine to provide to the installation program. The key will be used for SSH authentication onto your cluster nodes for debugging and disaster recovery.
- You have obtained the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create an installation directory to store your required installation assets in:

```
$ mkdir <installation_directory>
```



IMPORTANT

You must create a directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

2. Customize the sample **install-config.yaml** file template that is provided and save it in the **<installation_directory>**.



NOTE

You must name this configuration file **install-config.yaml**.

3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the next step of the installation process. You must back it up now.

Additional resources

- [Installation configuration parameters for GCP](#)

9.9.7.1. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 9.20. Minimum resource requirements

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Bootstrap	RHCOS	4	16 GB	100 GB	300

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Control plane	RHCOS	4	16 GB	100 GB	300
Compute	RHCOS, RHEL 8.6 and later [3]	2	8 GB	100 GB	300

1. One vCPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio: (threads per core × cores) × sockets = vCPUs.
2. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes which require a 10 ms p99 fsync duration. Note that on many cloud platforms, storage size and IOPS scale together, so you might need to over-allocate storage volume to obtain sufficient performance.
3. As with all user-provisioned installations, if you choose to use RHEL compute machines in your cluster, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Use of RHEL 7 compute machines is deprecated and has been removed in OpenShift Container Platform 4.10 and later.



NOTE

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

9.9.7.2. Tested instance types for GCP

The following Google Cloud Platform instance types have been tested with OpenShift Container Platform.

■

Example 9.46. Machine series

- **A2**
- **A3**
- **C2**
- **C2D**
- **C3**
- **C3D**
- **E2**
- **M1**
- **N1**
- **N2**
- **N2D**
- **Tau T2D**

9.9.7.3. Tested instance types for GCP on 64-bit ARM infrastructures

The following Google Cloud Platform (GCP) 64-bit ARM instance types have been tested with OpenShift Container Platform.

Example 9.47. Machine series for 64-bit ARM machines

- **Tau T2A**

9.9.7.4. Using custom machine types

Using a custom machine type to install a OpenShift Container Platform cluster is supported.

Consider the following when using a custom machine type:

- Similar to predefined instance types, custom machine types must meet the minimum resource requirements for control plane and compute machines. For more information, see "Minimum resource requirements for cluster installation".
- The name of the custom machine type must adhere to the following syntax:
custom-<number_of_cpus>-<amount_of_memory_in_mb>

For example, **custom-6-20480**.

As part of the installation process, you specify the custom machine type in the **install-config.yaml** file.

Sample install-config.yaml file with a custom machine type


```

compute:
- architecture: amd64
  hyperthreading: Enabled
  name: worker
  platform:
    gcp:
      type: custom-6-20480
  replicas: 2
controlPlane:
  architecture: amd64
  hyperthreading: Enabled
  name: master
  platform:
    gcp:
      type: custom-6-20480
  replicas: 3

```

9.9.7.5. Enabling Shielded VMs

You can use Shielded VMs when installing your cluster. Shielded VMs have extra security features including secure boot, firmware and integrity monitoring, and rootkit detection. For more information, see Google's documentation on [Shielded VMs](#).



NOTE

Shielded VMs are currently not supported on clusters with 64-bit ARM infrastructures.

Prerequisites

- You have created an **install-config.yaml** file.

Procedure

- Use a text editor to edit the **install-config.yaml** file prior to deploying your cluster and add one of the following stanzas:
 - To use shielded VMs for only control plane machines:

```

controlPlane:
  platform:
    gcp:
      secureBoot: Enabled

```

- To use shielded VMs for only compute machines:

```

compute:
- platform:
    gcp:
      secureBoot: Enabled

```

- To use shielded VMs for all machines:

```

platform:
  gcp:

```

```
defaultMachinePlatform:
  secureBoot: Enabled
```

9.9.7.6. Enabling Confidential VMs

You can use Confidential VMs when installing your cluster. Confidential VMs encrypt data while it is being processed. For more information, see Google's documentation on [Confidential Computing](#). You can enable Confidential VMs and Shielded VMs at the same time, although they are not dependent on each other.



NOTE

Confidential VMs are currently not supported on 64-bit ARM architectures.

Prerequisites

- You have created an **install-config.yaml** file.

Procedure

- Use a text editor to edit the **install-config.yaml** file prior to deploying your cluster and add one of the following stanzas:
 - To use confidential VMs for only control plane machines:

```
controlPlane:
  platform:
    gcp:
      confidentialCompute: Enabled 1
      type: n2d-standard-8 2
      onHostMaintenance: Terminate 3
```

- Enable confidential VMs.
- Specify a machine type that supports Confidential VMs. Confidential VMs require the N2D or C2D series of machine types. For more information on supported machine types, see [Supported operating systems and machine types](#).
- Specify the behavior of the VM during a host maintenance event, such as a hardware or software update. For a machine that uses Confidential VM, this value must be set to **Terminate**, which stops the VM. Confidential VMs do not support live VM migration.

- To use confidential VMs for only compute machines:


```
compute:
- platform:
  gcp:
    confidentialCompute: Enabled
    type: n2d-standard-8
    onHostMaintenance: Terminate
```

- To use confidential VMs for all machines:

```
platform:
  gcp:
    defaultMachinePlatform:
      confidentialCompute: Enabled
      type: n2d-standard-8
      onHostMaintenance: Terminate
```

9.9.7.7. Sample customized install-config.yaml file for GCP

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.



IMPORTANT

This sample YAML file is provided for reference only. You must obtain your **install-config.yaml** file by using the installation program and modify it.

```
apiVersion: v1
baseDomain: example.com 1
credentialsMode: Mint 2
controlPlane: 3 4
  hyperthreading: Enabled 5
  name: master
  platform:
    gcp:
      type: n2-standard-4
      zones:
        - us-central1-a
        - us-central1-c
      osDisk:
        diskType: pd-ssd
        diskSizeGB: 1024
        encryptionKey: 6
          kmsKey:
            name: worker-key
            keyRing: test-machine-keys
            location: global
            projectID: project-id
      tags: 7
        - control-plane-tag1
        - control-plane-tag2
      osImage: 8
        project: example-project-name
        name: example-image-name
    replicas: 3
  compute: 9 10
    - hyperthreading: Enabled 11
    name: worker
    platform:
      gcp:
        type: n2-standard-4
        zones:
          - us-central1-a
```

```

- us-central1-c
osDisk:
  diskType: pd-standard
  diskSizeGB: 128
  encryptionKey: 12
  kmsKey:
    name: worker-key
    keyRing: test-machine-keys
    location: global
    projectID: project-id
  tags: 13
- compute-tag1
- compute-tag2
osImage: 14
  project: example-project-name
  name: example-image-name
replicas: 3
metadata:
  name: test-cluster 15
networking:
  clusterNetwork:
    - cidr: 10.128.0.0/14
    hostPrefix: 23
  machineNetwork:
    - cidr: 10.0.0.0/16
  networkType: OVNKubernetes 16
  serviceNetwork:
    - 172.30.0.0/16
platform:
  gcp:
    projectID: openshift-production 17
    region: us-central1 18
    defaultMachinePlatform:
      tags: 19
      - global-tag1
      - global-tag2
    osImage: 20
      project: example-project-name
      name: example-image-name
    network: existing_vpc 21
    controlPlaneSubnet: control_plane_subnet 22
    computeSubnet: compute_subnet 23
  pullSecret: '{"auths": ...}' 24
  fips: false 25
  sshKey: ssh-ed25519 AAAA... 26
  publish: Internal 27

```

1 15 17 18 24 Required. The installation program prompts you for this value.

2 Optional: Add this parameter to force the Cloud Credential Operator (CCO) to use the specified mode. By default, the CCO uses the root credentials in the **kube-system** namespace to dynamically try to determine the capabilities of the credentials. For details about CCO modes, see the "About the Cloud Credential Operator" section in the *Authentication and authorization* guide.

- 3 9 If you do not provide these parameters and values, the installation program provides the default value.
- 4 10 The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.
- 5 11 Whether to enable or disable simultaneous multithreading, or **hyperthreading**. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores. You can disable it by setting the parameter value to **Disabled**. If you disable simultaneous multithreading in some cluster machines, you must disable it in all cluster machines.



IMPORTANT

If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance. Use larger machine types, such as **n1-standard-8**, for your machines if you disable simultaneous multithreading.

- 6 12 Optional: The custom encryption key section to encrypt both virtual machines and persistent volumes. Your default compute service account must have the permissions granted to use your KMS key and have the correct IAM role assigned. The default service account name follows the **service-<project_number>@compute-system.iam.gserviceaccount.com** pattern. For more information about granting the correct permissions for your service account, see "Machine management" → "Creating compute machine sets" → "Creating a compute machine set on GCP".
- 7 13 19 Optional: A set of network tags to apply to the control plane or compute machine sets. The **platform.gcp.defaultMachinePlatform.tags** parameter will apply to both control plane and compute machines. If the **compute.platform.gcp.tags** or **controlPlane.platform.gcp.tags** parameters are set, they override the **platform.gcp.defaultMachinePlatform.tags** parameter.
- 8 14 20 Optional: A custom Red Hat Enterprise Linux CoreOS (RHCOS) that should be used to boot control plane and compute machines. The **project** and **name** parameters under **platform.gcp.defaultMachinePlatform.osImage** apply to both control plane and compute machines. If the **project** and **name** parameters under **controlPlane.platform.gcp.osImage** or **compute.platform.gcp.osImage** are set, they override the **platform.gcp.defaultMachinePlatform.osImage** parameters.
- 16 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 21 Specify the name of an existing VPC.
- 22 Specify the name of the existing subnet to deploy the control plane machines to. The subnet must belong to the VPC that you specified.
- 23 Specify the name of the existing subnet to deploy the compute machines to. The subnet must belong to the VPC that you specified.
- 25 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.




IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 26 You can optionally provide the **sshKey** value that you use to access the machines in your cluster.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- 27 How to publish the user-facing endpoints of your cluster. Set **publish** to **Internal** to deploy a private cluster, which cannot be accessed from the internet. The default value is **External**.

Additional resources

- [Enabling customer-managed encryption keys for a compute machine set](#)

9.9.7.8. Create an Ingress Controller with global access on GCP

You can create an Ingress Controller that has global access to a Google Cloud Platform (GCP) cluster. Global access is only available to Ingress Controllers using internal load balancers.

Prerequisites

- You created the **install-config.yaml** and complete any modifications to it.

Procedure

Create an Ingress Controller with global access on a new GCP cluster.

1. Change to the directory that contains the installation program and create a manifest file:

```
$ ./openshift-install create manifests --dir <installation_directory> 1
```

- 1 For **<installation_directory>**, specify the name of the directory that contains the **install-config.yaml** file for your cluster.

2. Create a file that is named **cluster-ingress-default-ingresscontroller.yaml** in the **<installation_directory>/manifests/** directory:

```
$ touch <installation_directory>/manifests/cluster-ingress-default-ingresscontroller.yaml 1
```

- 1 For **<installation_directory>**, specify the directory name that contains the **manifests/** directory for your cluster.

After creating the file, several network configuration files are in the **manifests/** directory, as shown:

```
$ ls <installation_directory>/manifests/cluster-ingress-default-ingresscontroller.yaml
```

Example output

```
cluster-ingress-default-ingresscontroller.yaml
```

3. Open the **cluster-ingress-default-ingresscontroller.yaml** file in an editor and enter a custom resource (CR) that describes the Operator configuration you want:

Sample **clientAccess** configuration to **Global**

```
apiVersion: operator.openshift.io/v1
kind: IngressController
metadata:
  name: default
  namespace: openshift-ingress-operator
spec:
  endpointPublishingStrategy:
    loadBalancer:
      providerParameters:
        gcp:
          clientAccess: Global 1
          type: GCP
          scope: Internal 2
          type: LoadBalancerService
```

- 1 Set **gcp.clientAccess** to **Global**.
- 2 Global access is only available to Ingress Controllers using internal load balancers.

9.9.7.9. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.


For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:


```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.



NOTE

The installation program does not support the proxy **readinessEndpoints** field.

**NOTE**

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.


The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.

**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

9.9.8. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.

**IMPORTANT**

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**. To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.



NOTE

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your PATH.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

9.9.9. Alternatives to storing administrator-level secrets in the kube-system project

By default, administrator secrets are stored in the **kube-system** project. If you configured the **credentialsMode** parameter in the **install-config.yaml** file to **Manual**, you must use one of the following alternatives:

- To manage long-term cloud credentials manually, follow the procedure in [Manually creating long-term credentials](#).
- To implement short-term credentials that are managed outside the cluster for individual components, follow the procedures in [Configuring a GCP cluster to use short-term credentials](#).

9.9.9.1. Manually creating long-term credentials

The Cloud Credential Operator (CCO) can be put into manual mode prior to installation in environments where the cloud identity and access management (IAM) APIs are not reachable, or the administrator prefers not to store an administrator-level credential secret in the cluster **kube-system** namespace.

Procedure

1. Add the following granular permissions to the GCP account that the installation program uses:

Example 9.48. Required GCP permissions

- compute.machineTypes.list
- compute.regions.list
- compute.zones.list
- dns.changes.create
- dns.changes.get
- dns.managedZones.create
- dns.managedZones.delete
- dns.managedZones.get
- dns.managedZones.list
- dns.networks.bindPrivateDNSZone
- dns.resourceRecordSets.create
- dns.resourceRecordSets.delete
- dns.resourceRecordSets.list

- If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```
apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...
```

- If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

- Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

- Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included \1 \
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \2 \
  --to=<path_to_directory_for_credentials_requests> \3
```

- The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- Specify the location of the **install-config.yaml** file.
- Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

This command creates a YAML file for each **CredentialsRequest** object.

Sample CredentialsRequest object

```
apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
```

```

kind: GCPProviderSpec
predefinedRoles:
- roles/storage.admin
- roles/iam.serviceAccountUser
skipServiceCheck: true
...

```

6. Create YAML files for secrets in the **openshift-install** manifests directory that you generated previously. The secrets must be stored using the namespace and secret name defined in the **spec.secretRef** for each **CredentialsRequest** object.

Sample CredentialsRequest object with secrets

```

apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  name: <component_credentials_request>
  namespace: openshift-cloud-credential-operator
  ...
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    ...
  secretRef:
    name: <component_secret>
    namespace: <component_namespace>
  ...


```

Sample Secret object

```

apiVersion: v1
kind: Secret
metadata:
  name: <component_secret>
  namespace: <component_namespace>
data:
  service_account.json: <base64_encoded_gcp_service_account_file>

```



IMPORTANT


Before upgrading a cluster that uses manually maintained credentials, you must ensure that the CCO is in an upgradeable state.

9.9.9.2. Configuring a GCP cluster to use short-term credentials

To install a cluster that is configured to use GCP Workload Identity, you must configure the CCO utility and create the required GCP resources for your cluster.

9.9.9.2.1. Configuring the Cloud Credential Operator utility

To create and manage cloud credentials from outside of the cluster when the Cloud Credential Operator (CCO) is operating in manual mode, extract and prepare the CCO utility (**ccoctl**) binary.

**NOTE**

The **ccoctl** utility is a Linux binary that must run in a Linux environment.

Prerequisites

- You have access to an OpenShift Container Platform account with cluster administrator access.
- You have installed the OpenShift CLI (**oc**).
- You have added one of the following authentication options to the GCP account that the installation program uses:
 - The **IAM Workload Identity Pool Admin** role.
 - The following granular permissions:

Example 9.49. Required GCP permissions

- `compute.projects.get`
- `iam.googleapis.com/workloadIdentityPoolProviders.create`
- `iam.googleapis.com/workloadIdentityPoolProviders.get`
- `iam.googleapis.com/workloadIdentityPools.create`
- `iam.googleapis.com/workloadIdentityPools.delete`
- `iam.googleapis.com/workloadIdentityPools.get`
- `iam.googleapis.com/workloadIdentityPools.undelete`
- `iam.roles.create`
- `iam.roles.delete`
- `iam.roles.list`
- `iam.roles.undelete`
- `iam.roles.update`
- `iam.serviceAccounts.create`
- `iam.serviceAccounts.delete`
- `iam.serviceAccounts.getIamPolicy`
- `iam.serviceAccounts.list`
- `iam.serviceAccounts.setIamPolicy`
- `iam.workloadIdentityPoolProviders.get`
- `iam.workloadIdentityPools.delete`
- `resourceManager.projects.get`

- `resourcemanager.projects.getIamPolicy`
- `resourcemanager.projects.setIamPolicy`
- `storage.buckets.create`
- `storage.buckets.delete`
- `storage.buckets.get`
- `storage.buckets.getIamPolicy`
- `storage.buckets.setIamPolicy`
- `storage.objects.create`
- `storage.objects.delete`
- `storage.objects.list`

Procedure

1. Set a variable for the OpenShift Container Platform release image by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

2. Obtain the CCO container image from the OpenShift Container Platform release image by running the following command:

```
$ CCO_IMAGE=$(oc adm release info --image-for='cloud-credential-operator'
$RELEASE_IMAGE -a ~/.pull-secret)
```



NOTE

Ensure that the architecture of the **\$RELEASE_IMAGE** matches the architecture of the environment in which you will use the **ccoctl** tool.

3. Extract the **ccoctl** binary from the CCO container image within the OpenShift Container Platform release image by running the following command:

```
$ oc image extract $CCO_IMAGE \
--file="/usr/bin/ccoctl.<rhel_version>" \
-a ~/.pull-secret
```

- 1 For **<rhel_version>**, specify the value that corresponds to the version of Red Hat Enterprise Linux (RHEL) that the host uses. If no value is specified, **ccoctl.rhel8** is used by default. The following values are valid:

- **rhel8**: Specify this value for hosts that use RHEL 8.
- **rhel9**: Specify this value for hosts that use RHEL 9.

- Change the permissions to make **ccoctl** executable by running the following command:

```
$ chmod 775 ccoctl.<rhel_version>
```

Verification

- To verify that **ccoctl** is ready to use, display the help file by running the following command:

```
$ ccoctl --help
```

Output of **ccoctl --help**

```
OpenShift credentials provisioning tool
```

```
Usage:
```

```
ccoctl [command]
```

```
Available Commands:
```

```
aws      Manage credentials objects for AWS cloud
azure    Manage credentials objects for Azure
gcp      Manage credentials objects for Google cloud
help     Help about any command
ibmcloud Manage credentials objects for IBM Cloud
nutanix  Manage credentials objects for Nutanix
```

```
Flags:
```

```
-h, --help  help for ccoctl
```

```
Use "ccoctl [command] --help" for more information about a command.
```

9.9.9.2.2. Creating GCP resources with the Cloud Credential Operator utility

You can use the **ccoctl gcp create-all** command to automate the creation of GCP resources.



NOTE

By default, **ccoctl** creates objects in the directory in which the commands are run. To create the objects in a different directory, use the **--output-dir** flag. This procedure uses **<path_to_ccoctl_output_dir>** to refer to this directory.

Prerequisites

You must have:

- Extracted and prepared the **ccoctl** binary.

Procedure


- Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```


2. Extract the list of **CredentialsRequest** objects from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included \1
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \2
  --to=<path_to_directory_for_credentials_requests> \3
```

- 1 The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2 Specify the location of the **install-config.yaml** file.
- 3 Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.




NOTE

This command might take a few moments to run.

3. Use the **ccoctl** tool to process all **CredentialsRequest** objects by running the following command:

```
$ ccoctl gcp create-all \
  --name=<name> \1
  --region=<gcp_region> \2
  --project=<gcp_project_id> \3
  --credentials-requests-dir=<path_to_credentials_requests_directory> \4
```

- 1 Specify the user-defined name for all created GCP resources used for tracking.
- 2 Specify the GCP region in which cloud resources will be created.
- 3 Specify the GCP project ID in which cloud resources will be created.
- 4 Specify the directory containing the files of **CredentialsRequest** manifests to create GCP service accounts.



NOTE

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

Verification

- To verify that the OpenShift Container Platform secrets are created, list the files in the **<path_to_ccoctl_output_dir>/manifests** directory:

```
$ ls <path_to_ccoctl_output_dir>/manifests
```

Example output

```
cluster-authentication-02-config.yaml
openshift-cloud-controller-manager-gcp-ccm-cloud-credentials-credentials.yaml
openshift-cloud-credential-operator-cloud-credential-operator-gcp-ro-creds-credentials.yaml
openshift-cloud-network-config-controller-cloud-credentials-credentials.yaml
openshift-cluster-api-capg-manager-bootstrap-credentials-credentials.yaml
openshift-cluster-csi-drivers-gcp-pd-cloud-credentials-credentials.yaml
openshift-image-registry-installer-cloud-credentials-credentials.yaml
openshift-ingress-operator-cloud-credentials-credentials.yaml
openshift-machine-api-gcp-cloud-credentials-credentials.yaml
```

You can verify that the IAM service accounts are created by querying GCP. For more information, refer to GCP documentation on listing IAM service accounts.

9.9.9.2.3. Incorporating the Cloud Credential Operator utility manifests

To implement short-term security credentials managed outside the cluster for individual components, you must move the manifest files that the Cloud Credential Operator utility (**ccoctl**) created to the correct directories for the installation program.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have configured the Cloud Credential Operator utility (**ccoctl**).
- You have created the cloud provider resources that are required for your cluster with the **ccoctl** utility.

Procedure

1. Add the following granular permissions to the GCP account that the installation program uses:

Example 9.50. Required GCP permissions

- compute.machineTypes.list
- compute.regions.list
- compute.zones.list
- dns.changes.create
- dns.changes.get
- dns.managedZones.create
- dns.managedZones.delete
- dns.managedZones.get
- dns.managedZones.list
- dns.networks.bindPrivateDNSZone

- `dns.resourceRecordSets.create`
- `dns.resourceRecordSets.delete`
- `dns.resourceRecordSets.list`

- If you did not set the **credentialsMode** parameter in the **install-config.yaml** configuration file to **Manual**, modify the value as shown:

Sample configuration file snippet

```
apiVersion: v1
baseDomain: example.com
credentialsMode: Manual
# ...
```

- If you have not previously created installation manifest files, do so by running the following command:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation_directory>** is the directory in which the installation program creates files.

- Copy the manifests that the **ccoctl** utility generated to the **manifests** directory that the installation program created by running the following command:


```
$ cp /<path_to_ccoctl_output_dir>/manifests/* ./manifests/
```

- Copy the private key that the **ccoctl** utility generated in the **tls** directory to the installation directory by running the following command:

```
$ cp -a /<path_to_ccoctl_output_dir>/tls .
```

9.9.10. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.

Procedure

1. Remove any existing GCP credentials that do not use the service account key for the GCP account that you configured for your cluster and that are stored in the following locations:
 - The **GOOGLE_CREDENTIALS**, **GOOGLE_CLOUD_KEYFILE_JSON**, or **GKLOUD_KEYFILE_JSON** environment variables
 - The `~/.gcp/osServiceAccount.json` file
 - The **gcloud cli** default credentials
2. Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

1 For `<installation_directory>`, specify the location of your customized `./install-config.yaml` file.

2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

3. Optional: You can reduce the number of permissions for the service account that you used to install the cluster.
 - If you assigned the **Owner** role to your service account, you can remove that role and replace it with the **Viewer** role.
 - If you included the **Service Account Key Admin** role, you can remove it.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to `<installation_directory>/.openshift_install.log`.



IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

9.9.11. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

Additional resources

- See [Accessing the web console](#) for more details about accessing and understanding the OpenShift Container Platform web console.

9.9.12. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service

9.9.13. Next steps

- [Customize your cluster](#).
- If necessary, you can [opt out of remote health reporting](#) .

9.10. INSTALLING A CLUSTER ON USER-PROVISIONED INFRASTRUCTURE IN GCP BY USING DEPLOYMENT MANAGER TEMPLATES

In OpenShift Container Platform version 4.16, you can install a cluster on Google Cloud Platform (GCP) that uses infrastructure that you provide.

The steps for performing a user-provided infrastructure install are outlined here. Several [Deployment Manager](#) templates are provided to assist in completing these steps or to help model your own. You are also free to create the required resources through other methods.



IMPORTANT

The steps for performing a user-provisioned infrastructure installation are provided as an example only. Installing a cluster with infrastructure you provide requires knowledge of the cloud provider and the installation process of OpenShift Container Platform. Several Deployment Manager templates are provided to assist in completing these steps or to help model your own. You are also free to create the required resources through other methods; the templates are just an example.

9.10.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- If you use a firewall and plan to use the Telemetry service, you [configured the firewall to allow the sites](#) that your cluster requires access to.
- If the cloud identity and access management (IAM) APIs are not accessible in your environment, or if you do not want to store an administrator-level credential secret in the **kube-system** namespace, you can [manually create and maintain long-term credentials](#).

**NOTE**

Be sure to also review this site list if you are configuring a proxy.

9.10.2. Certificate signing requests management


Because your cluster has limited access to automatic machine management when you use infrastructure that you provision, you must provide a mechanism for approving cluster certificate signing requests (CSRs) after installation. The **kube-controller-manager** only approves the kubelet client CSRs. The **machine-approver** cannot guarantee the validity of a serving certificate that is requested by using kubelet credentials because it cannot confirm that the correct machine issued the request. You must determine and implement a method of verifying the validity of the kubelet serving certificate requests and approving them.

9.10.3. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.

**IMPORTANT**

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

9.10.4. Configuring your GCP project


Before you can install OpenShift Container Platform, you must configure a Google Cloud Platform (GCP) project to host it.

9.10.4.1. Creating a GCP project

To install OpenShift Container Platform, you must create a project in your Google Cloud Platform (GCP) account to host the cluster.

Procedure

- Create a project to host your OpenShift Container Platform cluster. See [Creating and Managing Projects](#) in the GCP documentation.



IMPORTANT

Your GCP project must use the Premium Network Service Tier if you are using installer-provisioned infrastructure. The Standard Network Service Tier is not supported for clusters installed using the installation program. The installation program configures internal load balancing for the **api-int.<cluster_name>.<base_domain>** URL; the Premium Tier is required for internal load balancing.

9.10.4.2. Enabling API services in GCP

Your Google Cloud Platform (GCP) project requires access to several API services to complete OpenShift Container Platform installation.

Prerequisites

- You created a project to host your cluster.

Procedure

- Enable the following required API services in the project that hosts your cluster. You may also enable optional API services which are not required for installation. See [Enabling services](#) in the GCP documentation.

Table 9.21. Required API services

API service	Console service name
Compute Engine API	compute.googleapis.com
Cloud Resource Manager API	cloudresourcemanager.googleapis.com
Google DNS API	dns.googleapis.com
IAM Service Account Credentials API	iamcredentials.googleapis.com
Identity and Access Management (IAM) API	iam.googleapis.com
Service Usage API	serviceusage.googleapis.com

Table 9.22. Optional API services

API service	Console service name
Cloud Deployment Manager V2 API	deploymentmanager.googleapis.com
Google Cloud APIs	cloudapis.googleapis.com
Service Management API	servicemanagement.googleapis.com


API service	Console service name
Google Cloud Storage JSON API	storage-api.googleapis.com
Cloud Storage	storage-component.googleapis.com

9.10.4.3. Configuring DNS for GCP

To install OpenShift Container Platform, the Google Cloud Platform (GCP) account you use must have a dedicated public hosted zone in the same project that you host the OpenShift Container Platform cluster. This zone must be authoritative for the domain. The DNS service provides cluster DNS resolution and name lookup for external connections to the cluster.

Procedure

1. Identify your domain, or subdomain, and registrar. You can transfer an existing domain and registrar or obtain a new one through GCP or another source.



NOTE

If you purchase a new domain, it can take time for the relevant DNS changes to propagate. For more information about purchasing domains through Google, see [Google Domains](#).

2. Create a public hosted zone for your domain or subdomain in your GCP project. See [Creating public zones](#) in the GCP documentation. Use an appropriate root domain, such as **openshiftcorp.com**, or subdomain, such as **clusters.openshiftcorp.com**.
3. Extract the new authoritative name servers from the hosted zone records. See [Look up your Cloud DNS name servers](#) in the GCP documentation. You typically have four name servers.
4. Update the registrar records for the name servers that your domain uses. For example, if you registered your domain to Google Domains, see the following topic in the Google Domains Help: [How to switch to custom name servers](#).
5. If you migrated your root domain to Google Cloud DNS, migrate your DNS records. See [Migrating to Cloud DNS](#) in the GCP documentation.
6. If you use a subdomain, follow your company's procedures to add its delegation records to the parent domain. This process might include a request to your company's IT department or the division that controls the root domain and DNS services for your company.


9.10.4.4. GCP account limits

The OpenShift Container Platform cluster uses a number of Google Cloud Platform (GCP) components, but the default [Quotas](#) do not affect your ability to install a default OpenShift Container Platform cluster.

A default cluster, which contains three compute and three control plane machines, uses the following resources. Note that some resources are required only during the bootstrap process and are removed after the cluster deploys.

Table 9.23. GCP resources used in a default cluster

Service	Component	Location	Total resources required	Resources removed after bootstrap
Service account	IAM	Global	6	1
Firewall rules	Networking	Global	11	1
Forwarding rules	Compute	Global	2	0
Health checks	Compute	Global	2	0
Images	Compute	Global	1	0
Networks	Networking	Global	1	0
Routers	Networking	Global	1	0
Routes	Networking	Global	2	0
Subnetworks	Compute	Global	2	0
Target pools	Networking	Global	2	0

**NOTE**

If any of the quotas are insufficient during installation, the installation program displays an error that states both which quota was exceeded and the region.

Be sure to consider your actual cluster size, planned cluster growth, and any usage from other clusters that are associated with your account. The CPU, static IP addresses, and persistent disk SSD (storage) quotas are the ones that are most likely to be insufficient.

If you plan to deploy your cluster in one of the following regions, you will exceed the maximum storage quota and are likely to exceed the CPU quota limit:

- **asia-east2**
- **asia-northeast2**
- **asia-south1**
- **australia-southeast1**
- **europa-north1**
- **europa-west2**
- **europa-west3**

- **europa-west6**
- **northamerica-northeast1**
- **southamerica-east1**
- **us-west2**

You can increase resource quotas from the [GCP console](#), but you might need to file a support ticket. Be sure to plan your cluster size early so that you can allow time to resolve the support ticket before you install your OpenShift Container Platform cluster.

9.10.4.5. Creating a service account in GCP

OpenShift Container Platform requires a Google Cloud Platform (GCP) service account that provides authentication and authorization to access data in the Google APIs. If you do not have an existing IAM service account that contains the required roles in your project, you must create one.

Prerequisites

- You created a project to host your cluster.

Procedure

1. Create a service account in the project that you use to host your OpenShift Container Platform cluster. See [Creating a service account](#) in the GCP documentation.
2. Grant the service account the appropriate permissions. You can either grant the individual permissions that follow or assign the **Owner** role to it. See [Granting roles to a service account for specific resources](#).




NOTE

While making the service account an owner of the project is the easiest way to gain the required permissions, it means that service account has complete control over the project. You must determine if the risk that comes from offering that power is acceptable.

3. You can create the service account key in JSON format, or attach the service account to a GCP virtual machine. See [Creating service account keys](#) and [Creating and enabling service accounts for instances](#) in the GCP documentation.

You must have a service account key or a virtual machine with an attached service account to create the cluster.



NOTE

If you use a virtual machine with an attached service account to create your cluster, you must set **credentialsMode: Manual** in the **install-config.yaml** file before installation.

9.10.4.6. Required GCP roles

When you attach the **Owner** role to the service account that you create, you grant that service account all permissions, including those that are required to install OpenShift Container Platform. If your

organization's security policies require a more restrictive set of permissions, you can create a service account with the following permissions. If you deploy your cluster into an existing virtual private cloud (VPC), the service account does not require certain networking permissions, which are noted in the following lists:

Required roles for the installation program

- Compute Admin
- Role Administrator
- Security Admin
- Service Account Admin
- Service Account Key Admin
- Service Account User
- Storage Admin

Required roles for creating network resources during installation

- DNS Administrator

Required roles for using the Cloud Credential Operator in passthrough mode

- Compute Load Balancer Admin

Required roles for user-provisioned GCP infrastructure

- Deployment Manager Editor

The following roles are applied to the service accounts that the control plane and compute machines use:

Table 9.24. GCP service account roles

Account	Roles
Control Plane	roles/compute.instanceAdmin
	roles/compute.networkAdmin
	roles/compute.securityAdmin
	roles/storage.admin
	roles/iam.serviceAccountUser
Compute	roles/compute.viewer
	roles/storage.admin

Account	Roles

9.10.4.7. Required GCP permissions for user-provisioned infrastructure

When you attach the **Owner** role to the service account that you create, you grant that service account all permissions, including those that are required to install OpenShift Container Platform.

If your organization's security policies require a more restrictive set of permissions, you can create [custom roles](#) with the necessary permissions. The following permissions are required for the user-provisioned infrastructure for creating and deleting the OpenShift Container Platform cluster.

Example 9.51. Required permissions for creating network resources

- **compute.addresses.create**
- **compute.addresses.createInternal**
- **compute.addresses.delete**
- **compute.addresses.get**
- **compute.addresses.list**
- **compute.addresses.use**
- **compute.addresses.useInternal**
- **compute.firewalls.create**
- **compute.firewalls.delete**
- **compute.firewalls.get**
- **compute.firewalls.list**
- **compute.forwardingRules.create**
- **compute.forwardingRules.get**
- **compute.forwardingRules.list**
- **compute.forwardingRules.setLabels**
- **compute.networks.create**
- **compute.networks.get**
- **compute.networks.list**
- **compute.networks.updatePolicy**
- **compute.routers.create**
- **compute.routers.get**

- **compute.routers.list**
- **compute.routers.update**
- **compute.routes.list**
- **compute.subnetworks.create**
- **compute.subnetworks.get**
- **compute.subnetworks.list**
- **compute.subnetworks.use**
- **compute.subnetworks.useExternallp**

Example 9.52. Required permissions for creating load balancer resources

- **compute.regionBackendServices.create**
- **compute.regionBackendServices.get**
- **compute.regionBackendServices.list**
- **compute.regionBackendServices.update**
- **compute.regionBackendServices.use**
- **compute.targetPools.addInstance**
- **compute.targetPools.create**
- **compute.targetPools.get**
- **compute.targetPools.list**
- **compute.targetPools.removeInstance**
- **compute.targetPools.use**

Example 9.53. Required permissions for creating DNS resources

- **dns.changes.create**
- **dns.changes.get**
- **dns.managedZones.create**
- **dns.managedZones.get**
- **dns.managedZones.list**
- **dns.networks.bindPrivateDNSZone**
- **dns.resourceRecordSets.create**

- `dns.resourceRecordSets.list`
- `dns.resourceRecordSets.update`

Example 9.54. Required permissions for creating Service Account resources

- `iam.serviceAccountKeys.create`
- `iam.serviceAccountKeys.delete`
- `iam.serviceAccountKeys.get`
- `iam.serviceAccountKeys.list`
- `iam.serviceAccounts.actAs`
- `iam.serviceAccounts.create`
- `iam.serviceAccounts.delete`
- `iam.serviceAccounts.get`
- `iam.serviceAccounts.list`
- `resourcemanager.projects.get`
- `resourcemanager.projects.getIamPolicy`
- `resourcemanager.projects.setIamPolicy`

Example 9.55. Required permissions for creating compute resources

- `compute.disks.create`
- `compute.disks.get`
- `compute.disks.list`
- `compute.instanceGroups.create`
- `compute.instanceGroups.delete`
- `compute.instanceGroups.get`
- `compute.instanceGroups.list`
- `compute.instanceGroups.update`
- `compute.instanceGroups.use`
- `compute.instances.create`
- `compute.instances.delete`
- `compute.instances.get`

- **compute.instances.list**
- **compute.instances.setLabels**
- **compute.instances.setMetadata**
- **compute.instances.setServiceAccount**
- **compute.instances.setTags**
- **compute.instances.use**
- **compute.machineTypes.get**
- **compute.machineTypes.list**

Example 9.56. Required for creating storage resources

- **storage.buckets.create**
- **storage.buckets.delete**
- **storage.buckets.get**
- **storage.buckets.list**
- **storage.objects.create**
- **storage.objects.delete**
- **storage.objects.get**
- **storage.objects.list**

Example 9.57. Required permissions for creating health check resources

- **compute.healthChecks.create**
- **compute.healthChecks.get**
- **compute.healthChecks.list**
- **compute.healthChecks.useReadOnly**
- **compute.httpHealthChecks.create**
- **compute.httpHealthChecks.get**
- **compute.httpHealthChecks.list**
- **compute.httpHealthChecks.useReadOnly**

Example 9.58. Required permissions to get GCP zone and region related information

- `compute.globalOperations.get`
- `compute.regionOperations.get`
- `compute.regions.list`
- `compute.zoneOperations.get`
- `compute.zones.get`
- `compute.zones.list`

Example 9.59. Required permissions for checking services and quotas

- `monitoring.timeSeries.list`
- `serviceusage.quotas.get`
- `serviceusage.services.list`

Example 9.60. Required IAM permissions for installation

- `iam.roles.get`

Example 9.61. Required Images permissions for installation

- `compute.images.create`
- `compute.images.delete`
- `compute.images.get`
- `compute.images.list`

Example 9.62. Optional permission for running gather bootstrap

- `compute.instances.getSerialPortOutput`

Example 9.63. Required permissions for deleting network resources

- `compute.addresses.delete`
- `compute.addresses.deleteInternal`
- `compute.addresses.list`
- `compute.firewalls.delete`
- `compute.firewalls.list`

- `compute.forwardingRules.delete`
- `compute.forwardingRules.list`
- `compute.networks.delete`
- `compute.networks.list`
- `compute.networks.updatePolicy`
- `compute.routers.delete`
- `compute.routers.list`
- `compute.routes.list`
- `compute.subnetworks.delete`
- `compute.subnetworks.list`

Example 9.64. Required permissions for deleting load balancer resources

- `compute.regionBackendServices.delete`
- `compute.regionBackendServices.list`
- `compute.targetPools.delete`
- `compute.targetPools.list`

Example 9.65. Required permissions for deleting DNS resources

- `dns.changes.create`
- `dns.managedZones.delete`
- `dns.managedZones.get`
- `dns.managedZones.list`
- `dns.resourceRecordSets.delete`
- `dns.resourceRecordSets.list`

Example 9.66. Required permissions for deleting Service Account resources

- `iam.serviceAccounts.delete`
- `iam.serviceAccounts.get`
- `iam.serviceAccounts.list`
- `resourcemanager.projects.getIamPolicy`

- `resourceManager.projects.setIamPolicy`

Example 9.67. Required permissions for deleting compute resources

- `compute.disks.delete`
- `compute.disks.list`
- `compute.instanceGroups.delete`
- `compute.instanceGroups.list`
- `compute.instances.delete`
- `compute.instances.list`
- `compute.instances.stop`
- `compute.machineTypes.list`

Example 9.68. Required for deleting storage resources

- `storage.buckets.delete`
- `storage.buckets.getIamPolicy`
- `storage.buckets.list`
- `storage.objects.delete`
- `storage.objects.list`

Example 9.69. Required permissions for deleting health check resources

- `compute.healthChecks.delete`
- `compute.healthChecks.list`
- `compute.httpHealthChecks.delete`
- `compute.httpHealthChecks.list`

Example 9.70. Required Images permissions for deletion

- `compute.images.delete`
- `compute.images.list`

Example 9.71. Required permissions to get Region related information

- **compute.regions.get**

Example 9.72. Required Deployment Manager permissions

- **deploymentmanager.deployments.create**
- **deploymentmanager.deployments.delete**
- **deploymentmanager.deployments.get**
- **deploymentmanager.deployments.list**
- **deploymentmanager.manifests.get**
- **deploymentmanager.operations.get**
- **deploymentmanager.resources.list**

Additional resources

- [Optimizing storage](#)

9.10.4.8. Supported GCP regions

You can deploy an OpenShift Container Platform cluster to the following Google Cloud Platform (GCP) regions:

- **africa-south1** (Johannesburg, South Africa)
- **asia-east1** (Changhua County, Taiwan)
- **asia-east2** (Hong Kong)
- **asia-northeast1** (Tokyo, Japan)
- **asia-northeast2** (Osaka, Japan)
- **asia-northeast3** (Seoul, South Korea)
- **asia-south1** (Mumbai, India)
- **asia-south2** (Delhi, India)
- **asia-southeast1** (Jurong West, Singapore)
- **asia-southeast2** (Jakarta, Indonesia)
- **australia-southeast1** (Sydney, Australia)
- **australia-southeast2** (Melbourne, Australia)
- **eu-central2** (Warsaw, Poland)
- **eu-north1** (Hamina, Finland)

- **europa-southwest1** (Madrid, Spain)
- **europa-west1** (St. Ghislain, Belgium)
- **europa-west2** (London, England, UK)
- **europa-west3** (Frankfurt, Germany)
- **europa-west4** (Eemshaven, Netherlands)
- **europa-west6** (Zürich, Switzerland)
- **europa-west8** (Milan, Italy)
- **europa-west9** (Paris, France)
- **europa-west12** (Turin, Italy)
- **me-central1** (Doha, Qatar, Middle East)
- **me-central2** (Dammam, Saudi Arabia, Middle East)
- **me-west1** (Tel Aviv, Israel)
- **northamerica-northeast1** (Montréal, Québec, Canada)
- **northamerica-northeast2** (Toronto, Ontario, Canada)
- **southamerica-east1** (São Paulo, Brazil)
- **southamerica-west1** (Santiago, Chile)
- **us-central1** (Council Bluffs, Iowa, USA)
- **us-east1** (Moncks Corner, South Carolina, USA)
- **us-east4** (Ashburn, Northern Virginia, USA)
- **us-east5** (Columbus, Ohio)
- **us-south1** (Dallas, Texas)
- **us-west1** (The Dalles, Oregon, USA)
- **us-west2** (Los Angeles, California, USA)
- **us-west3** (Salt Lake City, Utah, USA)
- **us-west4** (Las Vegas, Nevada, USA)



NOTE

To determine which machine type instances are available by region and zone, see the Google [documentation](#).

9.10.4.9. Installing and configuring CLI tools for GCP

To install OpenShift Container Platform on Google Cloud Platform (GCP) using user-provisioned infrastructure, you must install and configure the CLI tools for GCP.

Prerequisites

- You created a project to host your cluster.
- You created a service account and granted it the required permissions.

Procedure

1. Install the following binaries in **\$PATH**:

- **gcloud**
- **gsutil**

See [Install the latest Cloud SDK version](#) in the GCP documentation.

2. Authenticate using the **gcloud** tool with your configured service account.
See [Authorizing with a service account](#) in the GCP documentation.

9.10.5. Requirements for a cluster with user-provisioned infrastructure

For a cluster that contains user-provisioned infrastructure, you must deploy all of the required machines.

This section describes the requirements for deploying OpenShift Container Platform on user-provisioned infrastructure.

9.10.5.1. Required machines for cluster installation

The smallest OpenShift Container Platform clusters require the following hosts:

Table 9.25. Minimum required hosts

Hosts	Description
One temporary bootstrap machine	The cluster requires the bootstrap machine to deploy the OpenShift Container Platform cluster on the three control plane machines. You can remove the bootstrap machine after you install the cluster.
Three control plane machines	The control plane machines run the Kubernetes and OpenShift Container Platform services that form the control plane.
At least two compute machines, which are also known as worker machines.	The workloads requested by OpenShift Container Platform users run on the compute machines.



IMPORTANT

To maintain high availability of your cluster, use separate physical hosts for these cluster machines.

The bootstrap and control plane machines must use Red Hat Enterprise Linux CoreOS (RHCOS) as the operating system. However, the compute machines can choose between Red Hat Enterprise Linux CoreOS (RHCOS), Red Hat Enterprise Linux (RHEL) 8.6 and later.

Note that RHCOS is based on Red Hat Enterprise Linux (RHEL) 9.2 and inherits all of its hardware certifications and requirements. See [Red Hat Enterprise Linux technology capabilities and limits](#).

9.10.5.2. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 9.26. Minimum resource requirements

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Bootstrap	RHCOS	4	16 GB	100 GB	300
Control plane	RHCOS	4	16 GB	100 GB	300
Compute	RHCOS, RHEL 8.6 and later [3]	2	8 GB	100 GB	300

1. One vCPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio: $(\text{threads per core} \times \text{cores}) \times \text{sockets} = \text{vCPUs}$.
2. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes which require a 10 ms p99 fsync duration. Note that on many cloud platforms, storage size and IOPS scale together, so you might need to over-allocate storage volume to obtain sufficient performance.
3. As with all user-provisioned installations, if you choose to use RHEL compute machines in your cluster, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Use of RHEL 7 compute machines is deprecated and has been removed in OpenShift Container Platform 4.10 and later.

**NOTE**

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

9.10.5.3. Tested instance types for GCP

The following Google Cloud Platform instance types have been tested with OpenShift Container Platform.

Example 9.73. Machine series

- **A2**
- **A3**
- **C2**
- **C2D**
- **C3**
- **C3D**
- **E2**
- **M1**
- **N1**
- **N2**
- **N2D**
- **Tau T2D**

9.10.5.4. Tested instance types for GCP on 64-bit ARM infrastructures

The following Google Cloud Platform (GCP) 64-bit ARM instance types have been tested with OpenShift Container Platform.

Example 9.74. Machine series for 64-bit ARM machines

- **Tau T2A**

9.10.5.5. Using custom machine types

Using a custom machine type to install a OpenShift Container Platform cluster is supported.

Consider the following when using a custom machine type:

- Similar to predefined instance types, custom machine types must meet the minimum resource requirements for control plane and compute machines. For more information, see "Minimum resource requirements for cluster installation".
- The name of the custom machine type must adhere to the following syntax:
custom-`<number_of_cpus>`-`<amount_of_memory_in_mb>`

For example, **custom-6-20480**.

9.10.6. Creating the installation files for GCP

To install OpenShift Container Platform on Google Cloud Platform (GCP) using user-provisioned infrastructure, you must generate the files that the installation program needs to deploy your cluster and modify them so that the cluster creates only the machines that it will use. You generate and customize the **install-config.yaml** file, Kubernetes manifests, and Ignition config files. You also have the option to first set up a separate **var** partition during the preparation phases of installation.

9.10.6.1. Optional: Creating a separate `/var` partition

It is recommended that disk partitioning for OpenShift Container Platform be left to the installer. However, there are cases where you might want to create separate partitions in a part of the filesystem that you expect to grow.


OpenShift Container Platform supports the addition of a single partition to attach storage to either the **/var** partition or a subdirectory of **/var**. For example:

- **/var/lib/containers**: Holds container-related content that can grow as more images and containers are added to a system.
- **/var/lib/etcd**: Holds data that you might want to keep separate for purposes such as performance optimization of etcd storage.
- **/var**: Holds data that you might want to keep separate for purposes such as auditing.

Storing the contents of a **/var** directory separately makes it easier to grow storage for those areas as needed and reinstall OpenShift Container Platform at a later date and keep that data intact. With this method, you will not have to pull all your containers again, nor will you have to copy massive log files when you update systems.

Because **/var** must be in place before a fresh installation of Red Hat Enterprise Linux CoreOS (RHCOS), the following procedure sets up the separate **/var** partition by creating a machine config manifest that is

inserted during the **openshift-install** preparation phases of an OpenShift Container Platform installation.



IMPORTANT

If you follow the steps to create a separate **/var** partition in this procedure, it is not necessary to create the Kubernetes manifest and Ignition config files again as described later in this section.

Procedure

1. Create a directory to hold the OpenShift Container Platform installation files:

```
$ mkdir $HOME/clusterconfig
```

2. Run **openshift-install** to create a set of files in the **manifest** and **openshift** subdirectories. Answer the system questions as you are prompted:

```
$ openshift-install create manifests --dir $HOME/clusterconfig
```

Example output

```
? SSH Public Key ...
INFO Credentials loaded from the "myprofile" profile in file "/home/myuser/.aws/credentials"
INFO Consuming Install Config from target directory
INFO Manifests created in: $HOME/clusterconfig/manifests and
$HOME/clusterconfig/openshift
```

3. Optional: Confirm that the installation program created manifests in the **clusterconfig/openshift** directory:

```
$ ls $HOME/clusterconfig/openshift/
```

Example output

```
99_kubeadmin-password-secret.yaml
99_openshift-cluster-api_master-machines-0.yaml
99_openshift-cluster-api_master-machines-1.yaml
99_openshift-cluster-api_master-machines-2.yaml
...
```

4. Create a Butane config that configures the additional partition. For example, name the file **\$HOME/clusterconfig/98-var-partition.bu**, change the disk device name to the name of the storage device on the **worker** systems, and set the storage size as appropriate. This example places the **/var** directory on a separate partition:

```
variant: openshift
version: 4.16.0
metadata:
  labels:
    machineconfiguration.openshift.io/role: worker
  name: 98-var-partition
storage:
```

```

disks:
- device: /dev/disk/by-id/<device_name> ❶
  partitions:
  - label: var
    start_mib: <partition_start_offset> ❷
    size_mib: <partition_size> ❸
    number: 5
  filesystems:
  - device: /dev/disk/by-partlabel/var
    path: /var
    format: xfs
    mount_options: [defaults, prjquota] ❹
    with_mount_unit: true

```

- ❶ The storage device name of the disk that you want to partition.
- ❷ When adding a data partition to the boot disk, a minimum value of 25000 MiB (Mebibytes) is recommended. The root file system is automatically resized to fill all available space up to the specified offset. If no value is specified, or if the specified value is smaller than the recommended minimum, the resulting root file system will be too small, and future reinstalls of RHCOS might overwrite the beginning of the data partition.
- ❸ The size of the data partition in mebibytes.
- ❹ The **prjquota** mount option must be enabled for filesystems used for container storage.



NOTE

When creating a separate **/var** partition, you cannot use different instance types for worker nodes, if the different instance types do not have the same device name.

5. Create a manifest from the Butane config and save it to the **clusterconfig/openshift** directory. For example, run the following command:

```
$ butane $HOME/clusterconfig/98-var-partition.bu -o $HOME/clusterconfig/openshift/98-var-partition.yaml
```

6. Run **openshift-install** again to create Ignition configs from a set of files in the **manifest** and **openshift** subdirectories:

```
$ openshift-install create ignition-configs --dir $HOME/clusterconfig
$ ls $HOME/clusterconfig/
auth bootstrap.ign master.ign metadata.json worker.ign
```

Now you can use the Ignition config files as input to the installation procedures to install Red Hat Enterprise Linux CoreOS (RHCOS) systems.

9.10.6.2. Creating the installation configuration file

You can customize the OpenShift Container Platform cluster you install on Google Cloud Platform (GCP).

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create the **install-config.yaml** file.
 - a. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install create install-config --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the directory name to store the files that the installation program creates.

When specifying the directory:

- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
 - Use an empty directory. Some installation assets, such as bootstrap X.509 certificates, have short expiration intervals, therefore you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.
- b. At the prompts, provide the configuration details for your cloud:
 - i. Optional: Select an SSH key to use to access your cluster machines.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- ii. Select **gcp** as the platform to target.
- iii. If you have not configured the service account key for your GCP account on your computer, you must obtain it from GCP and paste the contents of the file or enter the absolute path to the file.
- iv. Select the project ID to provision the cluster in. The default value is specified by the service account that you configured.
- v. Select the region to deploy the cluster to.
- vi. Select the base domain to deploy the cluster to. The base domain corresponds to the public DNS zone that you created for your cluster.
- vii. Enter a descriptive name for your cluster.

2. Modify the **install-config.yaml** file. You can find more information about the available parameters in the "Installation configuration parameters" section.



NOTE

If you are installing a three-node cluster, be sure to set the **compute.replicas** parameter to **0**. This ensures that the cluster's control planes are schedulable. For more information, see "Installing a three-node cluster on GCP".

3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

Additional resources

- [Installation configuration parameters for GCP](#)

9.10.6.3. Enabling Shielded VMs

You can use Shielded VMs when installing your cluster. Shielded VMs have extra security features including secure boot, firmware and integrity monitoring, and rootkit detection. For more information, see Google's documentation on [Shielded VMs](#).



NOTE

Shielded VMs are currently not supported on clusters with 64-bit ARM infrastructures.

Prerequisites

- You have created an **install-config.yaml** file.

Procedure

- Use a text editor to edit the **install-config.yaml** file prior to deploying your cluster and add one of the following stanzas:
 - a. To use shielded VMs for only control plane machines:

```
controlPlane:
  platform:
    gcp:
      secureBoot: Enabled
```

- b. To use shielded VMs for only compute machines:

```
compute:
- platform:
  gcp:
    secureBoot: Enabled
```

- c. To use shielded VMs for all machines:

```
platform:
  gcp:
    defaultMachinePlatform:
      secureBoot: Enabled
```

9.10.6.4. Enabling Confidential VMs

You can use Confidential VMs when installing your cluster. Confidential VMs encrypt data while it is being processed. For more information, see Google's documentation on [Confidential Computing](#). You can enable Confidential VMs and Shielded VMs at the same time, although they are not dependent on each other.



NOTE

Confidential VMs are currently not supported on 64-bit ARM architectures.

Prerequisites

- You have created an **install-config.yaml** file.

Procedure

- Use a text editor to edit the **install-config.yaml** file prior to deploying your cluster and add one of the following stanzas:
 - To use confidential VMs for only control plane machines:

```
controlPlane:
  platform:
    gcp:
      confidentialCompute: Enabled 1
      type: n2d-standard-8 2
      onHostMaintenance: Terminate 3
```

- Enable confidential VMs.
- Specify a machine type that supports Confidential VMs. Confidential VMs require the N2D or C2D series of machine types. For more information on supported machine types, see [Supported operating systems and machine types](#).
- Specify the behavior of the VM during a host maintenance event, such as a hardware or software update. For a machine that uses Confidential VM, this value must be set to **Terminate**, which stops the VM. Confidential VMs do not support live VM migration.

- To use confidential VMs for only compute machines:

```
compute:
  - platform:
      gcp:
```

```
confidentialCompute: Enabled
type: n2d-standard-8
onHostMaintenance: Terminate
```

- c. To use confidential VMs for all machines:

```
platform:
  gcp:
    defaultMachinePlatform:
      confidentialCompute: Enabled
      type: n2d-standard-8
      onHostMaintenance: Terminate
```

9.10.6.5. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.


For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.

**NOTE**

The installation program does not support the proxy **readinessEndpoints** field.


**NOTE**

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.


**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

9.10.6.6. Creating the Kubernetes manifest and Ignition config files

Because you must modify some cluster definition files and manually start the cluster machines, you must generate the Kubernetes manifest and Ignition config files that the cluster needs to configure the machines.

The installation configuration file transforms into the Kubernetes manifests. The manifests wrap into the Ignition configuration files, which are later used to configure the cluster machines.



IMPORTANT

- The Ignition config files that the OpenShift Container Platform installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

Prerequisites

- You obtained the OpenShift Container Platform installation program.
- You created the **install-config.yaml** installation configuration file.

Procedure

1. Change to the directory that contains the OpenShift Container Platform installation program and generate the Kubernetes manifests for the cluster:

```
$ ./openshift-install create manifests --dir <installation_directory> 1
```

- 1 For **<installation_directory>**, specify the installation directory that contains the **install-config.yaml** file you created.

2. Remove the Kubernetes manifest files that define the control plane machines:

```
$ rm -f <installation_directory>/openshift/99_openshift-cluster-api_master-machines-*.yaml
```


By removing these files, you prevent the cluster from automatically generating control plane machines.

3. Remove the Kubernetes manifest files that define the control plane machine set:

```
$ rm -f <installation_directory>/openshift/99_openshift-machine-api_master-control-plane-machine-set.yaml
```

4. Optional: If you do not want the cluster to provision compute machines, remove the Kubernetes manifest files that define the worker machines:

```
$ rm -f <installation_directory>/openshift/99_openshift-cluster-api_worker-machineset-*.yaml
```


**IMPORTANT**

If you disabled the **MachineAPI** capability when installing a cluster on user-provisioned infrastructure, you must remove the Kubernetes manifest files that define the worker machines. Otherwise, your cluster fails to install.

Because you create and manage the worker machines yourself, you do not need to initialize these machines.

**WARNING**

If you are installing a three-node cluster, skip the following step to allow the control plane nodes to be schedulable.

**IMPORTANT**

When you configure control plane nodes from the default unschedulable to schedulable, additional subscriptions are required. This is because control plane nodes then become compute nodes.

5. Check that the **mastersSchedulable** parameter in the `<installation_directory>/manifests/cluster-scheduler-02-config.yml` Kubernetes manifest file is set to **false**. This setting prevents pods from being scheduled on the control plane machines:
 - a. Open the `<installation_directory>/manifests/cluster-scheduler-02-config.yml` file.
 - b. Locate the **mastersSchedulable** parameter and ensure that it is set to **false**.
 - c. Save and exit the file.
6. Optional: If you do not want [the Ingress Operator](#) to create DNS records on your behalf, remove the **privateZone** and **publicZone** sections from the `<installation_directory>/manifests/cluster-dns-02-config.yml` DNS configuration file:

```
apiVersion: config.openshift.io/v1
kind: DNS
metadata:
  creationTimestamp: null
  name: cluster
spec:
  baseDomain: example.openshift.com
  privateZone: ❶
    id: mycluster-100419-private-zone
  publicZone: ❷
    id: example.openshift.com
status: {}
```

- ❶ ❷ Remove this section completely.


If you do so, you must add ingress DNS records manually in a later step.

- To create the Ignition configuration files, run the following command from the directory that contains the installation program:

```
$ ./openshift-install create ignition-configs --dir <installation_directory> 1
```

- For **<installation_directory>**, specify the same installation directory.

Ignition config files are created for the bootstrap, control plane, and compute nodes in the installation directory. The **kubeadmin-password** and **kubeconfig** files are created in the **./<installation_directory>/auth** directory:



Additional resources

- [Optional: Adding the ingress DNS records](#)

9.10.7. Exporting common variables

9.10.7.1. Extracting the infrastructure name

The Ignition config files contain a unique cluster identifier that you can use to uniquely identify your cluster in Google Cloud Platform (GCP). The infrastructure name is also used to locate the appropriate GCP resources during an OpenShift Container Platform installation. The provided Deployment Manager templates contain references to this infrastructure name, so you must extract it.

Prerequisites

- You obtained the OpenShift Container Platform installation program and the pull secret for your cluster.
- You generated the Ignition config files for your cluster.
- You installed the **jq** package.

Procedure

- To extract and view the infrastructure name from the Ignition config file metadata, run the following command:

```
$ jq -r .infraID <installation_directory>/metadata.json 1
```

- For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

Example output

```
openshift-vw9j6 1
```

- 1** The output of this command is your cluster name and a random string.

9.10.7.2. Exporting common variables for Deployment Manager templates

You must export a common set of variables that are used with the provided Deployment Manager templates used to assist in completing a user-provided infrastructure install on Google Cloud Platform (GCP).



NOTE

Specific Deployment Manager templates can also require additional exported variables, which are detailed in their related procedures.

Prerequisites

- Obtain the OpenShift Container Platform installation program and the pull secret for your cluster.
- Generate the Ignition config files for your cluster.
- Install the **jq** package.

Procedure

1. Export the following common variables to be used by the provided Deployment Manager templates:


```
$ export BASE_DOMAIN='<base_domain>'
$ export BASE_DOMAIN_ZONE_NAME='<base_domain_zone_name>'
$ export NETWORK_CIDR='10.0.0.0/16'
$ export MASTER_SUBNET_CIDR='10.0.0.0/17'
$ export WORKER_SUBNET_CIDR='10.0.128.0/17'

$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
$ export CLUSTER_NAME=`jq -r .clusterName <installation_directory>/metadata.json`
$ export INFRA_ID=`jq -r .infraID <installation_directory>/metadata.json`
$ export PROJECT_NAME=`jq -r .gcp.projectID <installation_directory>/metadata.json`
$ export REGION=`jq -r .gcp.region <installation_directory>/metadata.json`
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

9.10.8. Creating a VPC in GCP

You must create a VPC in Google Cloud Platform (GCP) for your OpenShift Container Platform cluster to use. You can customize the VPC to meet your requirements. One way to create the VPC is to modify the provided Deployment Manager template.



NOTE

If you do not use the provided Deployment Manager template to create your GCP infrastructure, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- Configure a GCP account.
- Generate the Ignition config files for your cluster.

Procedure

1. Copy the template from the **Deployment Manager template for the VPC** section of this topic and save it as **01_vpc.py** on your computer. This template describes the VPC that your cluster requires.
2. Create a **01_vpc.yaml** resource definition file:

```
$ cat <<EOF >01_vpc.yaml
imports:
- path: 01_vpc.py

resources:
- name: cluster-vpc
  type: 01_vpc.py
  properties:
    infra_id: '${INFRA_ID}' 1
    region: '${REGION}' 2
    master_subnet_cidr: '${MASTER_SUBNET_CIDR}' 3
    worker_subnet_cidr: '${WORKER_SUBNET_CIDR}' 4
EOF
```

- 1** **infra_id** is the **INFRA_ID** infrastructure name from the extraction step.
- 2** **region** is the region to deploy the cluster into, for example **us-central1**.
- 3** **master_subnet_cidr** is the CIDR for the master subnet, for example **10.0.0.0/17**.
- 4** **worker_subnet_cidr** is the CIDR for the worker subnet, for example **10.0.128.0/17**.

3. Create the deployment by using the **gcloud** CLI:

```
$ gcloud deployment-manager deployments create ${INFRA_ID}-vpc --config 01_vpc.yaml
```

9.10.8.1. Deployment Manager template for the VPC

You can use the following Deployment Manager template to deploy the VPC that you need for your OpenShift Container Platform cluster:

Example 9.75. 01_vpc.py Deployment Manager template

-

```

def GenerateConfig(context):

resources = [{
  'name': context.properties['infra_id'] + '-network',
  'type': 'compute.v1.network',
  'properties': {
    'region': context.properties['region'],
    'autoCreateSubnetworks': False
  }
}, {
  'name': context.properties['infra_id'] + '-master-subnet',
  'type': 'compute.v1.subnetwork',
  'properties': {
    'region': context.properties['region'],
    'network': '${ref.' + context.properties['infra_id'] + '-network.selfLink}',
    'ipCidrRange': context.properties['master_subnet_cidr']
  }
}, {
  'name': context.properties['infra_id'] + '-worker-subnet',
  'type': 'compute.v1.subnetwork',
  'properties': {
    'region': context.properties['region'],
    'network': '${ref.' + context.properties['infra_id'] + '-network.selfLink}',
    'ipCidrRange': context.properties['worker_subnet_cidr']
  }
}, {
  'name': context.properties['infra_id'] + '-router',
  'type': 'compute.v1.router',
  'properties': {
    'region': context.properties['region'],
    'network': '${ref.' + context.properties['infra_id'] + '-network.selfLink}',
    'nats': [{
      'name': context.properties['infra_id'] + '-nat-master',
      'natIpAllocateOption': 'AUTO_ONLY',
      'minPortsPerVm': 7168,
      'sourceSubnetworkIpRangesToNat': 'LIST_OF_SUBNETWORKS',
      'subnetworks': [{
        'name': '${ref.' + context.properties['infra_id'] + '-master-subnet.selfLink}',
        'sourceIpRangesToNat': ['ALL_IP_RANGES']
      }
    ]
  }
}, {
  'name': context.properties['infra_id'] + '-nat-worker',
  'natIpAllocateOption': 'AUTO_ONLY',
  'minPortsPerVm': 512,
  'sourceSubnetworkIpRangesToNat': 'LIST_OF_SUBNETWORKS',
  'subnetworks': [{
    'name': '${ref.' + context.properties['infra_id'] + '-worker-subnet.selfLink}',
    'sourceIpRangesToNat': ['ALL_IP_RANGES']
  }
  ]
}
]}

return {'resources': resources}

```

9.10.9. Networking requirements for user-provisioned infrastructure

All the Red Hat Enterprise Linux CoreOS (RHCOS) machines require networking to be configured in **initramfs** during boot to fetch their Ignition config files.

9.10.9.1. Setting the cluster node hostnames through DHCP

On Red Hat Enterprise Linux CoreOS (RHCOS) machines, the hostname is set through NetworkManager. By default, the machines obtain their hostname through DHCP. If the hostname is not provided by DHCP, set statically through kernel arguments, or another method, it is obtained through a reverse DNS lookup. Reverse DNS lookup occurs after the network has been initialized on a node and can take time to resolve. Other system services can start prior to this and detect the hostname as **localhost** or similar. You can avoid this by using DHCP to provide the hostname for each cluster node.

Additionally, setting the hostnames through DHCP can bypass any manual DNS record name configuration errors in environments that have a DNS split-horizon implementation.

9.10.9.2. Network connectivity requirements

You must configure the network connectivity between machines to allow OpenShift Container Platform cluster components to communicate. Each machine must be able to resolve the hostnames of all other machines in the cluster.

This section provides details about the ports that are required.



IMPORTANT

In connected OpenShift Container Platform environments, all nodes are required to have internet access to pull images for platform containers and provide telemetry data to Red Hat.

Table 9.27. Ports used for all-machine to all-machine communications

Protocol	Port	Description
ICMP	N/A	Network reachability tests
TCP	1936	Metrics
	9000-9999	Host level services, including the node exporter on ports 9100-9101 and the Cluster Version Operator on port 9099 .
	10250-10259	The default ports that Kubernetes reserves
UDP	4789	VXLAN
	6081	Geneve
	9000-9999	Host level services, including the node exporter on ports 9100-9101 .

Protocol	Port	Description
	500	IPsec IKE packets
	4500	IPsec NAT-T packets
	123	Network Time Protocol (NTP) on UDP port 123 If an external NTP time server is configured, you must open UDP port 123 .
TCP/UDP	30000-32767	Kubernetes node port
ESP	N/A	IPsec Encapsulating Security Payload (ESP)

Table 9.28. Ports used for all-machine to control plane communications

Protocol	Port	Description
TCP	6443	Kubernetes API

Table 9.29. Ports used for control plane machine to control plane machine communications

Protocol	Port	Description
TCP	2379-2380	etcd server and peer ports

9.10.10. Creating load balancers in GCP

You must configure load balancers in Google Cloud Platform (GCP) for your OpenShift Container Platform cluster to use. One way to create these components is to modify the provided Deployment Manager template.



NOTE

If you do not use the provided Deployment Manager template to create your GCP infrastructure, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- Configure a GCP account.
- Generate the Ignition config files for your cluster.
- Create and configure a VPC and associated subnets in GCP.

Procedure

1. Copy the template from the **Deployment Manager template for the internal load balancer** section of this topic and save it as **02_lb_int.py** on your computer. This template describes the internal load balancing objects that your cluster requires.
2. For an external cluster, also copy the template from the **Deployment Manager template for the external load balancer** section of this topic and save it as **02_lb_ext.py** on your computer. This template describes the external load balancing objects that your cluster requires.
3. Export the variables that the deployment template uses:

- a. Export the cluster network location:

```
$ export CLUSTER_NETWORK=(`gcloud compute networks describe ${INFRA_ID}-
network --format json | jq -r .selfLink`)
```

- b. Export the control plane subnet location:

```
$ export CONTROL_SUBNET=(`gcloud compute networks subnets describe
${INFRA_ID}-master-subnet --region=${REGION} --format json | jq -r .selfLink`)
```

- c. Export the three zones that the cluster uses:

```
$ export ZONE_0=(`gcloud compute regions describe ${REGION} --format=json | jq -r
.zones[0] | cut -d "/" -f9`)
```

```
$ export ZONE_1=(`gcloud compute regions describe ${REGION} --format=json | jq -r
.zones[1] | cut -d "/" -f9`)
```

```
$ export ZONE_2=(`gcloud compute regions describe ${REGION} --format=json | jq -r
.zones[2] | cut -d "/" -f9`)
```

4. Create a **02_infra.yaml** resource definition file:

```
$ cat <<EOF >02_infra.yaml
imports:
- path: 02_lb_ext.py
- path: 02_lb_int.py ❶
resources:
- name: cluster-lb-ext ❷
  type: 02_lb_ext.py
  properties:
    infra_id: '${INFRA_ID}' ❸
    region: '${REGION}' ❹
- name: cluster-lb-int
  type: 02_lb_int.py
  properties:
    cluster_network: '${CLUSTER_NETWORK}'
    control_subnet: '${CONTROL_SUBNET}' ❺
    infra_id: '${INFRA_ID}'
    region: '${REGION}'
    zones: ❻
    - '${ZONE_0}'
```

```

- '${ZONE_1}'
- '${ZONE_2}'
EOF

```

1 **2** Required only when deploying an external cluster.

3 **infra_id** is the **INFRA_ID** infrastructure name from the extraction step.

4 **region** is the region to deploy the cluster into, for example **us-central1**.

5 **control_subnet** is the URI to the control subnet.

6 **zones** are the zones to deploy the control plane instances into, like **us-east1-b**, **us-east1-c**, and **us-east1-d**.

5. Create the deployment by using the **gcloud** CLI:

```
$ gcloud deployment-manager deployments create ${INFRA_ID}-infra --config 02_infra.yaml
```

6. Export the cluster IP address:

```
$ export CLUSTER_IP=`gcloud compute addresses describe ${INFRA_ID}-cluster-ip --
region=${REGION} --format json | jq -r .address`
```

7. For an external cluster, also export the cluster public IP address:

```
$ export CLUSTER_PUBLIC_IP=`gcloud compute addresses describe ${INFRA_ID}-cluster-
public-ip --region=${REGION} --format json | jq -r .address`
```

9.10.10.1. Deployment Manager template for the external load balancer

You can use the following Deployment Manager template to deploy the external load balancer that you need for your OpenShift Container Platform cluster:

Example 9.76. 02_lb_ext.py Deployment Manager template

```

def GenerateConfig(context):

    resources = [{
        'name': context.properties['infra_id'] + '-cluster-public-ip',
        'type': 'compute.v1.address',
        'properties': {
            'region': context.properties['region']
        }
    }, {
        # Refer to docs/dev/kube-apiserver-health-check.md on how to correctly setup health check
        # probe for kube-apiserver
        'name': context.properties['infra_id'] + '-api-http-health-check',
        'type': 'compute.v1.httpHealthCheck',
        'properties': {
            'port': 6080,
            'requestPath': '/readyz'
        }
    }

```

```

    }, {
      'name': context.properties['infra_id'] + '-api-target-pool',
      'type': 'compute.v1.targetPool',
      'properties': {
        'region': context.properties['region'],
        'healthChecks': ['$$(ref.' + context.properties['infra_id'] + '-api-http-health-check.selfLink)'],
        'instances': []
      }
    }, {
      'name': context.properties['infra_id'] + '-api-forwarding-rule',
      'type': 'compute.v1.forwardingRule',
      'properties': {
        'region': context.properties['region'],
        'IPAddress': '$(ref.' + context.properties['infra_id'] + '-cluster-public-ip.selfLink)',
        'target': '$(ref.' + context.properties['infra_id'] + '-api-target-pool.selfLink)',
        'portRange': '6443'
      }
    }
  ]
}

return {'resources': resources}

```

9.10.10.2. Deployment Manager template for the internal load balancer

You can use the following Deployment Manager template to deploy the internal load balancer that you need for your OpenShift Container Platform cluster:

Example 9.77. `02_lb_int.py` Deployment Manager template

```

def GenerateConfig(context):

    backends = []
    for zone in context.properties['zones']:
        backends.append({
            'group': '$(ref.' + context.properties['infra_id'] + '-master-' + zone + '-ig' + '.selfLink)'
        })

    resources = [{
        'name': context.properties['infra_id'] + '-cluster-ip',
        'type': 'compute.v1.address',
        'properties': {
            'addressType': 'INTERNAL',
            'region': context.properties['region'],
            'subnetwork': context.properties['control_subnet']
        }
    }, {
        # Refer to docs/dev/kube-apiserver-health-check.md on how to correctly setup health check
        # probe for kube-apiserver
        'name': context.properties['infra_id'] + '-api-internal-health-check',
        'type': 'compute.v1.healthCheck',
        'properties': {
            'httpsHealthCheck': {
                'port': 6443,
                'requestPath': '/readyz'
            }
        },
    },

```

```

        'type': "HTTPS"
    }
}, {
    'name': context.properties['infra_id'] + '-api-internal-backend-service',
    'type': 'compute.v1.regionBackendService',
    'properties': {
        'backends': backends,
        'healthChecks': ['$(ref.' + context.properties['infra_id'] + '-api-internal-health-
check.selfLink)'],
        'loadBalancingScheme': 'INTERNAL',
        'region': context.properties['region'],
        'protocol': 'TCP',
        'timeoutSec': 120
    }
}, {
    'name': context.properties['infra_id'] + '-api-internal-forwarding-rule',
    'type': 'compute.v1.forwardingRule',
    'properties': {
        'backendService': '$(ref.' + context.properties['infra_id'] + '-api-internal-backend-
service.selfLink)',
        'IPAddress': '$(ref.' + context.properties['infra_id'] + '-cluster-ip.selfLink)',
        'loadBalancingScheme': 'INTERNAL',
        'ports': ['6443','22623'],
        'region': context.properties['region'],
        'subnetwork': context.properties['control_subnet']
    }
}
]]

for zone in context.properties['zones']:
    resources.append({
        'name': context.properties['infra_id'] + '-master-' + zone + '-ig',
        'type': 'compute.v1.instanceGroup',
        'properties': {
            'namedPorts': [
                {
                    'name': 'ignition',
                    'port': 22623
                }, {
                    'name': 'https',
                    'port': 6443
                }
            ],
            'network': context.properties['cluster_network'],
            'zone': zone
        }
    })


return {'resources': resources}

```

You will need this template in addition to the **02_lb_ext.py** template when you create an external cluster.

9.10.11. Creating a private DNS zone in GCP

You must configure a private DNS zone in Google Cloud Platform (GCP) for your OpenShift Container Platform cluster to use. One way to create this component is to modify the provided Deployment Manager template.



NOTE

If you do not use the provided Deployment Manager template to create your GCP infrastructure, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- Configure a GCP account.
- Generate the Ignition config files for your cluster.
- Create and configure a VPC and associated subnets in GCP.

Procedure

1. Copy the template from the **Deployment Manager template for the private DNS** section of this topic and save it as **02_dns.py** on your computer. This template describes the private DNS objects that your cluster requires.
2. Create a **02_dns.yaml** resource definition file:

```
$ cat <<EOF >02_dns.yaml
imports:
- path: 02_dns.py

resources:
- name: cluster-dns
  type: 02_dns.py
  properties:
    infra_id: '${INFRA_ID}' ❶
    cluster_domain: '${CLUSTER_NAME}.${BASE_DOMAIN}' ❷
    cluster_network: '${CLUSTER_NETWORK}' ❸
EOF
```

- ❶ **infra_id** is the **INFRA_ID** infrastructure name from the extraction step.
- ❷ **cluster_domain** is the domain for the cluster, for example **openshift.example.com**.
- ❸ **cluster_network** is the **selfLink** URL to the cluster network.

3. Create the deployment by using the **gcloud** CLI:

```
$ gcloud deployment-manager deployments create ${INFRA_ID}-dns --config 02_dns.yaml
```

4. The templates do not create DNS entries due to limitations of Deployment Manager, so you must create them manually:
 - a. Add the internal DNS entries:

```

$ if [ -f transaction.yaml ]; then rm transaction.yaml; fi
$ gcloud dns record-sets transaction start --zone ${INFRA_ID}-private-zone
$ gcloud dns record-sets transaction add ${CLUSTER_IP} --name
api.${CLUSTER_NAME}.${BASE_DOMAIN}. --ttl 60 --type A --zone ${INFRA_ID}-
private-zone
$ gcloud dns record-sets transaction add ${CLUSTER_IP} --name api-
int.${CLUSTER_NAME}.${BASE_DOMAIN}. --ttl 60 --type A --zone ${INFRA_ID}-
private-zone
$ gcloud dns record-sets transaction execute --zone ${INFRA_ID}-private-zone

```

- b. For an external cluster, also add the external DNS entries:

```

$ if [ -f transaction.yaml ]; then rm transaction.yaml; fi
$ gcloud dns record-sets transaction start --zone ${BASE_DOMAIN_ZONE_NAME}
$ gcloud dns record-sets transaction add ${CLUSTER_PUBLIC_IP} --name
api.${CLUSTER_NAME}.${BASE_DOMAIN}. --ttl 60 --type A --zone
${BASE_DOMAIN_ZONE_NAME}
$ gcloud dns record-sets transaction execute --zone ${BASE_DOMAIN_ZONE_NAME}

```

9.10.11.1. Deployment Manager template for the private DNS

You can use the following Deployment Manager template to deploy the private DNS that you need for your OpenShift Container Platform cluster:

Example 9.78. `02_dns.py` Deployment Manager template

```

def GenerateConfig(context):

    resources = [{
        'name': context.properties['infra_id'] + '-private-zone',
        'type': 'dns.v1.managedZone',
        'properties': {
            'description': '',
            'dnsName': context.properties['cluster_domain'] + '.',
            'visibility': 'private',
            'privateVisibilityConfig': {
                'networks': [{
                    'networkUrl': context.properties['cluster_network']
                }]
            }
        }
    }]

    return {'resources': resources}

```

9.10.12. Creating firewall rules in GCP

You must create firewall rules in Google Cloud Platform (GCP) for your OpenShift Container Platform cluster to use. One way to create these components is to modify the provided Deployment Manager template.



NOTE

If you do not use the provided Deployment Manager template to create your GCP infrastructure, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- Configure a GCP account.
- Generate the Ignition config files for your cluster.
- Create and configure a VPC and associated subnets in GCP.

Procedure

1. Copy the template from the **Deployment Manager template for firewall rules** section of this topic and save it as **03_firewall.py** on your computer. This template describes the security groups that your cluster requires.
2. Create a **03_firewall.yaml** resource definition file:

```
$ cat <<EOF >03_firewall.yaml
imports:
- path: 03_firewall.py

resources:
- name: cluster-firewall
  type: 03_firewall.py
  properties:
    allowed_external_cidr: '0.0.0.0/0' 1
    infra_id: '${INFRA_ID}' 2
    cluster_network: '${CLUSTER_NETWORK}' 3
    network_cidr: '${NETWORK_CIDR}' 4
EOF
```

- 1 **allowed_external_cidr** is the CIDR range that can access the cluster API and SSH to the bootstrap host. For an internal cluster, set this value to **\${NETWORK_CIDR}**.
- 2 **infra_id** is the **INFRA_ID** infrastructure name from the extraction step.
- 3 **cluster_network** is the **selfLink** URL to the cluster network.
- 4 **network_cidr** is the CIDR of the VPC network, for example **10.0.0.0/16**.

3. Create the deployment by using the **gcloud** CLI:

```
$ gcloud deployment-manager deployments create ${INFRA_ID}-firewall --config
03_firewall.yaml
```

9.10.12.1. Deployment Manager template for firewall rules

You can use the following Deployment Manager template to deploy the firewall rules that you need for your OpenShift Container Platform cluster:

Example 9.79. 03_firewall.py Deployment Manager template

```
def GenerateConfig(context):

    resources = [{
        'name': context.properties['infra_id'] + '-bootstrap-in-ssh',
        'type': 'compute.v1.firewall',
        'properties': {
            'network': context.properties['cluster_network'],
            'allowed': [{
                'IPProtocol': 'tcp',
                'ports': ['22']
            }],
            'sourceRanges': [context.properties['allowed_external_cidr']],
            'targetTags': [context.properties['infra_id'] + '-bootstrap']
        }
    ], {
        'name': context.properties['infra_id'] + '-api',
        'type': 'compute.v1.firewall',
        'properties': {
            'network': context.properties['cluster_network'],
            'allowed': [{
                'IPProtocol': 'tcp',
                'ports': ['6443']
            }],
            'sourceRanges': [context.properties['allowed_external_cidr']],
            'targetTags': [context.properties['infra_id'] + '-master']
        }
    }, {
        'name': context.properties['infra_id'] + '-health-checks',
        'type': 'compute.v1.firewall',
        'properties': {
            'network': context.properties['cluster_network'],
            'allowed': [{
                'IPProtocol': 'tcp',
                'ports': ['6080', '6443', '22624']
            }],
            'sourceRanges': ['35.191.0.0/16', '130.211.0.0/22', '209.85.152.0/22', '209.85.204.0/22'],
            'targetTags': [context.properties['infra_id'] + '-master']
        }
    }, {
        'name': context.properties['infra_id'] + '-etcd',
        'type': 'compute.v1.firewall',
        'properties': {
            'network': context.properties['cluster_network'],
            'allowed': [{
                'IPProtocol': 'tcp',
                'ports': ['2379-2380']
            }],
            'sourceTags': [context.properties['infra_id'] + '-master'],
            'targetTags': [context.properties['infra_id'] + '-master']
        }
    }, {
```



```

'name': context.properties['infra_id'] + '-control-plane',
'type': 'compute.v1.firewall',
'properties': {
  'network': context.properties['cluster_network'],
  'allowed': [{
    'IPProtocol': 'tcp',
    'ports': ['10257']
  },{
    'IPProtocol': 'tcp',
    'ports': ['10259']
  },{
    'IPProtocol': 'tcp',
    'ports': ['22623']
  }],
  'sourceTags': [
    context.properties['infra_id'] + '-master',
    context.properties['infra_id'] + '-worker'
  ],
  'targetTags': [context.properties['infra_id'] + '-master']
}
}, {
'name': context.properties['infra_id'] + '-internal-network',
'type': 'compute.v1.firewall',
'properties': {
  'network': context.properties['cluster_network'],
  'allowed': [{
    'IPProtocol': 'icmp'
  },{
    'IPProtocol': 'tcp',
    'ports': ['22']
  }],
  'sourceRanges': [context.properties['network_cidr']],
  'targetTags': [
    context.properties['infra_id'] + '-master',
    context.properties['infra_id'] + '-worker'
  ]
}
}, {
'name': context.properties['infra_id'] + '-internal-cluster',
'type': 'compute.v1.firewall',
'properties': {
  'network': context.properties['cluster_network'],
  'allowed': [{
    'IPProtocol': 'udp',
    'ports': ['4789', '6081']
  },{
    'IPProtocol': 'udp',
    'ports': ['500', '4500']
  },{
    'IPProtocol': 'esp',
  },{
    'IPProtocol': 'tcp',
    'ports': ['9000-9999']
  },{
    'IPProtocol': 'udp',
    'ports': ['9000-9999']
  }
}
}

```

```

    },{
      'IPProtocol': 'tcp',
      'ports': ['10250']
    },{
      'IPProtocol': 'tcp',
      'ports': ['30000-32767']
    },{
      'IPProtocol': 'udp',
      'ports': ['30000-32767']
    }],
    'sourceTags': [
      context.properties['infra_id'] + '-master',
      context.properties['infra_id'] + '-worker'
    ],
    'targetTags': [
      context.properties['infra_id'] + '-master',
      context.properties['infra_id'] + '-worker'
    ]
  }
}
}
}

return {'resources': resources}

```

9.10.13. Creating IAM roles in GCP

You must create IAM roles in Google Cloud Platform (GCP) for your OpenShift Container Platform cluster to use. One way to create these components is to modify the provided Deployment Manager template.



NOTE

If you do not use the provided Deployment Manager template to create your GCP infrastructure, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- Configure a GCP account.
- Generate the Ignition config files for your cluster.
- Create and configure a VPC and associated subnets in GCP.

Procedure

1. Copy the template from the **Deployment Manager template for IAM roles** section of this topic and save it as **03_iam.py** on your computer. This template describes the IAM roles that your cluster requires.
2. Create a **03_iam.yaml** resource definition file:

```
$ cat <<EOF >03_iam.yaml
```

```

imports:
- path: 03_iam.py
resources:
- name: cluster-iam
  type: 03_iam.py
properties:
  infra_id: '${INFRA_ID}' ❶
EOF

```

- ❶ **infra_id** is the **INFRA_ID** infrastructure name from the extraction step.

3. Create the deployment by using the **gcloud** CLI:

```
$ gcloud deployment-manager deployments create ${INFRA_ID}-iam --config 03_iam.yaml
```

4. Export the variable for the master service account:

```
$ export MASTER_SERVICE_ACCOUNT=$(gcloud iam service-accounts list --filter
"email~^${INFRA_ID}-m@${PROJECT_NAME}." --format json | jq -r '[0].email')
```

5. Export the variable for the worker service account:

```
$ export WORKER_SERVICE_ACCOUNT=$(gcloud iam service-accounts list --filter
"email~^${INFRA_ID}-w@${PROJECT_NAME}." --format json | jq -r '[0].email')
```

6. Export the variable for the subnet that hosts the compute machines:

```
$ export COMPUTE_SUBNET=$(gcloud compute networks subnets describe ${INFRA_ID}-
worker-subnet --region=${REGION} --format json | jq -r .selfLink)
```

7. The templates do not create the policy bindings due to limitations of Deployment Manager, so you must create them manually:

```

$ gcloud projects add-iam-policy-binding ${PROJECT_NAME} --member
"serviceAccount:${MASTER_SERVICE_ACCOUNT}" --role "roles/compute.instanceAdmin"
$ gcloud projects add-iam-policy-binding ${PROJECT_NAME} --member
"serviceAccount:${MASTER_SERVICE_ACCOUNT}" --role "roles/compute.networkAdmin"
$ gcloud projects add-iam-policy-binding ${PROJECT_NAME} --member
"serviceAccount:${MASTER_SERVICE_ACCOUNT}" --role "roles/compute.securityAdmin"
$ gcloud projects add-iam-policy-binding ${PROJECT_NAME} --member
"serviceAccount:${MASTER_SERVICE_ACCOUNT}" --role "roles/iam.serviceAccountUser"
$ gcloud projects add-iam-policy-binding ${PROJECT_NAME} --member
"serviceAccount:${MASTER_SERVICE_ACCOUNT}" --role "roles/storage.admin"

$ gcloud projects add-iam-policy-binding ${PROJECT_NAME} --member
"serviceAccount:${WORKER_SERVICE_ACCOUNT}" --role "roles/compute.viewer"
$ gcloud projects add-iam-policy-binding ${PROJECT_NAME} --member
"serviceAccount:${WORKER_SERVICE_ACCOUNT}" --role "roles/storage.admin"

```

8. Create a service account key and store it locally for later use:

```
$ gcloud iam service-accounts keys create service-account-key.json --iam-
account=${MASTER_SERVICE_ACCOUNT}
```

-

9.10.13.1. Deployment Manager template for IAM roles

You can use the following Deployment Manager template to deploy the IAM roles that you need for your OpenShift Container Platform cluster:

Example 9.80. 03_iam.py Deployment Manager template

```
def GenerateConfig(context):

    resources = [{
        'name': context.properties['infra_id'] + '-master-node-sa',
        'type': 'iam.v1.serviceAccount',
        'properties': {
            'accountId': context.properties['infra_id'] + '-m',
            'displayName': context.properties['infra_id'] + '-master-node'
        }
    }, {
        'name': context.properties['infra_id'] + '-worker-node-sa',
        'type': 'iam.v1.serviceAccount',
        'properties': {
            'accountId': context.properties['infra_id'] + '-w',
            'displayName': context.properties['infra_id'] + '-worker-node'
        }
    }
    ]


    return {'resources': resources}
```

9.10.14. Creating the RHCOS cluster image for the GCP infrastructure

You must use a valid Red Hat Enterprise Linux CoreOS (RHCOS) image for Google Cloud Platform (GCP) for your OpenShift Container Platform nodes.

Procedure

1. Obtain the RHCOS image from the [RHCOS image mirror](#) page.



IMPORTANT

The RHCOS images might not change with every release of OpenShift Container Platform. You must download an image with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Use the image version that matches your OpenShift Container Platform version if it is available.

The file name contains the OpenShift Container Platform version number in the format **rhcos-
<version>-<arch>-gcp.<arch>.tar.gz**.

2. Create the Google storage bucket:

```
$ gsutil mb gs://<bucket_name>
```

3. Upload the RHCOS image to the Google storage bucket:

```
$ gsutil cp <downloaded_image_file_path>/rhcos-<version>-x86_64-gcp.x86_64.tar.gz
gs://<bucket_name>
```

4. Export the uploaded RHCOS image location as a variable:

```
$ export IMAGE_SOURCE=gs://<bucket_name>/rhcos-<version>-x86_64-gcp.x86_64.tar.gz
```

5. Create the cluster image:

```
$ gcloud compute images create "${INFRA_ID}-rhcos-image" \
--source-uri="${IMAGE_SOURCE}"
```

9.10.15. Creating the bootstrap machine in GCP

You must create the bootstrap machine in Google Cloud Platform (GCP) to use during OpenShift Container Platform cluster initialization. One way to create this machine is to modify the provided Deployment Manager template.



NOTE

If you do not use the provided Deployment Manager template to create your bootstrap machine, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- Configure a GCP account.
- Generate the Ignition config files for your cluster.
- Create and configure a VPC and associated subnets in GCP.
- Create and configure networking and load balancers in GCP.
- Create control plane and compute roles.
- Ensure pyOpenSSL is installed.

Procedure

1. Copy the template from the **Deployment Manager template for the bootstrap machine** section of this topic and save it as **04_bootstrap.py** on your computer. This template describes the bootstrap machine that your cluster requires.
2. Export the location of the Red Hat Enterprise Linux CoreOS (RHCOS) image that the installation program requires:

```
$ export CLUSTER_IMAGE=(`gcloud compute images describe ${INFRA_ID}-rhcos-image --
format json | jq -r .selfLink`)
```

3. Create a bucket and upload the **bootstrap.ign** file:

```
$ gsutil mb gs://${INFRA_ID}-bootstrap-ignition
```

```
$ gsutil cp <installation_directory>/bootstrap.ign gs://${INFRA_ID}-bootstrap-ignition/
```

4. Create a signed URL for the bootstrap instance to use to access the Ignition config. Export the URL from the output as a variable:

```
$ export BOOTSTRAP_IGN=`gsutil signurl -d 1h service-account-key.json gs://${INFRA_ID}-bootstrap-ignition/bootstrap.ign | grep "^gs:" | awk '{print $5}'`
```

5. Create a **04_bootstrap.yaml** resource definition file:

```
$ cat <<EOF >04_bootstrap.yaml
imports:
- path: 04_bootstrap.py

resources:
- name: cluster-bootstrap
  type: 04_bootstrap.py
  properties:
    infra_id: '${INFRA_ID}' 1
    region: '${REGION}' 2
    zone: '${ZONE_0}' 3

    cluster_network: '${CLUSTER_NETWORK}' 4
    control_subnet: '${CONTROL_SUBNET}' 5
    image: '${CLUSTER_IMAGE}' 6
    machine_type: 'n1-standard-4' 7
    root_volume_size: '128' 8

    bootstrap_ign: '${BOOTSTRAP_IGN}' 9
EOF
```

- 1 **infra_id** is the **INFRA_ID** infrastructure name from the extraction step.
- 2 **region** is the region to deploy the cluster into, for example **us-central1**.
- 3 **zone** is the zone to deploy the bootstrap instance into, for example **us-central1-b**.
- 4 **cluster_network** is the **selfLink** URL to the cluster network.
- 5 **control_subnet** is the **selfLink** URL to the control subnet.
- 6 **image** is the **selfLink** URL to the RHCOS image.
- 7 **machine_type** is the machine type of the instance, for example **n1-standard-4**.
- 8 **root_volume_size** is the boot disk size for the bootstrap machine.
- 9 **bootstrap_ign** is the URL output when creating a signed URL.

6. Create the deployment by using the **gcloud** CLI:

```
$ gcloud deployment-manager deployments create ${INFRA_ID}-bootstrap --config
04_bootstrap.yaml
```

7. The templates do not manage load balancer membership due to limitations of Deployment Manager, so you must add the bootstrap machine manually.

- a. Add the bootstrap instance to the internal load balancer instance group:

```
$ gcloud compute instance-groups unmanaged add-instances \
  ${INFRA_ID}-bootstrap-ig --zone=${ZONE_0} --instances=${INFRA_ID}-bootstrap
```

- b. Add the bootstrap instance group to the internal load balancer backend service:

```
$ gcloud compute backend-services add-backend \
  ${INFRA_ID}-api-internal-backend-service --region=${REGION} --instance-
group=${INFRA_ID}-bootstrap-ig --instance-group-zone=${ZONE_0}
```

9.10.15.1. Deployment Manager template for the bootstrap machine

You can use the following Deployment Manager template to deploy the bootstrap machine that you need for your OpenShift Container Platform cluster:

Example 9.81. 04_bootstrap.py Deployment Manager template

```
def GenerateConfig(context):

    resources = [{
        'name': context.properties['infra_id'] + '-bootstrap-public-ip',
        'type': 'compute.v1.address',
        'properties': {
            'region': context.properties['region']
        }
    }, {
        'name': context.properties['infra_id'] + '-bootstrap',
        'type': 'compute.v1.instance',
        'properties': {
            'disks': [{
                'autoDelete': True,
                'boot': True,
                'initializeParams': {
                    'diskSizeGb': context.properties['root_volume_size'],
                    'sourceImage': context.properties['image']
                }
            }],
            'machineType': 'zones/' + context.properties['zone'] + '/machineTypes/' +
context.properties['machine_type'],
            'metadata': {
                'items': [{
                    'key': 'user-data',
                    'value': '{"ignition":{"config":{"replace":{"source":"' + context.properties['bootstrap_ign']
+ '"},"version":"3.2.0"}}}',
                }],
            },
            'networkInterfaces': [{
```

```

        'subnetwork': context.properties['control_subnet'],
        'accessConfigs': [{
            'natIP': '${ref.' + context.properties['infra_id'] + '-bootstrap-public-ip.address}'
        }]
    }],
    'tags': {
        'items': [
            context.properties['infra_id'] + '-master',
            context.properties['infra_id'] + '-bootstrap'
        ]
    },
    'zone': context.properties['zone']
}
}, {
    'name': context.properties['infra_id'] + '-bootstrap-ig',
    'type': 'compute.v1.instanceGroup',
    'properties': {
        'namedPorts': [
            {
                'name': 'ignition',
                'port': 22623
            }, {
                'name': 'https',
                'port': 6443
            }
        ],
        'network': context.properties['cluster_network'],
        'zone': context.properties['zone']
    }
}
]]

return {'resources': resources}

```

9.10.16. Creating the control plane machines in GCP

You must create the control plane machines in Google Cloud Platform (GCP) for your cluster to use. One way to create these machines is to modify the provided Deployment Manager template.



NOTE

If you do not use the provided Deployment Manager template to create your control plane machines, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- Configure a GCP account.
- Generate the Ignition config files for your cluster.
- Create and configure a VPC and associated subnets in GCP.
- Create and configure networking and load balancers in GCP.

- Create control plane and compute roles.
- Create the bootstrap machine.

Procedure

1. Copy the template from the **Deployment Manager template for control plane machines** section of this topic and save it as **05_control_plane.py** on your computer. This template describes the control plane machines that your cluster requires.
2. Export the following variable required by the resource definition:

```
$ export MASTER_IGNITION=`cat <installation_directory>/master.ign`
```

3. Create a **05_control_plane.yaml** resource definition file:

```
$ cat <<EOF >05_control_plane.yaml
imports:
- path: 05_control_plane.py

resources:
- name: cluster-control-plane
  type: 05_control_plane.py
  properties:
    infra_id: '${INFRA_ID}' 1
    zones: 2
    - '${ZONE_0}'
    - '${ZONE_1}'
    - '${ZONE_2}'

    control_subnet: '${CONTROL_SUBNET}' 3
    image: '${CLUSTER_IMAGE}' 4
    machine_type: 'n1-standard-4' 5
    root_volume_size: '128'
    service_account_email: '${MASTER_SERVICE_ACCOUNT}' 6

    ignition: '${MASTER_IGNITION}' 7
EOF
```

- 1 **infra_id** is the **INFRA_ID** infrastructure name from the extraction step.
- 2 **zones** are the zones to deploy the control plane instances into, for example **us-central1-a**, **us-central1-b**, and **us-central1-c**.
- 3 **control_subnet** is the **selfLink** URL to the control subnet.
- 4 **image** is the **selfLink** URL to the RHCOS image.
- 5 **machine_type** is the machine type of the instance, for example **n1-standard-4**.
- 6 **service_account_email** is the email address for the master service account that you created.
- 7 **ignition** is the contents of the **master.ign** file.

4. Create the deployment by using the **gcloud** CLI:

```
$ gcloud deployment-manager deployments create ${INFRA_ID}-control-plane --config
05_control_plane.yaml
```

5. The templates do not manage load balancer membership due to limitations of Deployment Manager, so you must add the control plane machines manually.

- Run the following commands to add the control plane machines to the appropriate instance groups:

```
$ gcloud compute instance-groups unmanaged add-instances ${INFRA_ID}-master-
${ZONE_0}-ig --zone=${ZONE_0} --instances=${INFRA_ID}-master-0
```

```
$ gcloud compute instance-groups unmanaged add-instances ${INFRA_ID}-master-
${ZONE_1}-ig --zone=${ZONE_1} --instances=${INFRA_ID}-master-1
```

```
$ gcloud compute instance-groups unmanaged add-instances ${INFRA_ID}-master-
${ZONE_2}-ig --zone=${ZONE_2} --instances=${INFRA_ID}-master-2
```

- For an external cluster, you must also run the following commands to add the control plane machines to the target pools:

```
$ gcloud compute target-pools add-instances ${INFRA_ID}-api-target-pool --instances-
zone="${ZONE_0}" --instances=${INFRA_ID}-master-0
```

```
$ gcloud compute target-pools add-instances ${INFRA_ID}-api-target-pool --instances-
zone="${ZONE_1}" --instances=${INFRA_ID}-master-1
```

```
$ gcloud compute target-pools add-instances ${INFRA_ID}-api-target-pool --instances-
zone="${ZONE_2}" --instances=${INFRA_ID}-master-2
```

9.10.16.1. Deployment Manager template for control plane machines

You can use the following Deployment Manager template to deploy the control plane machines that you need for your OpenShift Container Platform cluster:

Example 9.82. 05_control_plane.py Deployment Manager template

```
def GenerateConfig(context):

    resources = [{
        'name': context.properties['infra_id'] + '-master-0',
        'type': 'compute.v1.instance',
        'properties': {
            'disks': [{
                'autoDelete': True,
                'boot': True,
                'initializeParams': {
                    'diskSizeGb': context.properties['root_volume_size'],
                    'diskType': 'zones/' + context.properties['zones'][0] + '/diskTypes/pd-ssd',
                    'sourceImage': context.properties['image']
```

```

    }
  }],
  'machineType': 'zones/' + context.properties['zones'][0] + '/machineTypes/' +
context.properties['machine_type'],
  'metadata': {
    'items': [{
      'key': 'user-data',
      'value': context.properties['ignition']
    }]
  },
  'networkInterfaces': [{
    'subnetwork': context.properties['control_subnet']
  }],
  'serviceAccounts': [{
    'email': context.properties['service_account_email'],
    'scopes': ['https://www.googleapis.com/auth/cloud-platform']
  }],
  'tags': {
    'items': [
      context.properties['infra_id'] + '-master',
    ]
  },
  'zone': context.properties['zones'][0]
}
}, {
  'name': context.properties['infra_id'] + '-master-1',
  'type': 'compute.v1.instance',
  'properties': {
    'disks': [{
      'autoDelete': True,
      'boot': True,
      'initializeParams': {
        'diskSizeGb': context.properties['root_volume_size'],
        'diskType': 'zones/' + context.properties['zones'][1] + '/diskTypes/pd-ssd',
        'sourceImage': context.properties['image']
      }
    }]
  },
  'machineType': 'zones/' + context.properties['zones'][1] + '/machineTypes/' +
context.properties['machine_type'],
  'metadata': {
    'items': [{
      'key': 'user-data',
      'value': context.properties['ignition']
    }]
  },
  'networkInterfaces': [{
    'subnetwork': context.properties['control_subnet']
  }],
  'serviceAccounts': [{
    'email': context.properties['service_account_email'],
    'scopes': ['https://www.googleapis.com/auth/cloud-platform']
  }],
  'tags': {
    'items': [
      context.properties['infra_id'] + '-master',
    ]
  }
}

```

```

    },
    'zone': context.properties['zones'][1]
  }
}, {
  'name': context.properties['infra_id'] + '-master-2',
  'type': 'compute.v1.instance',
  'properties': {
    'disks': [{
      'autoDelete': True,
      'boot': True,
      'initializeParams': {
        'diskSizeGb': context.properties['root_volume_size'],
        'diskType': 'zones/' + context.properties['zones'][2] + '/diskTypes/pd-ssd',
        'sourceImage': context.properties['image']
      }
    }
  ],
  'machineType': 'zones/' + context.properties['zones'][2] + '/machineTypes/' +
context.properties['machine_type'],
  'metadata': {
    'items': [{
      'key': 'user-data',
      'value': context.properties['ignition']
    }
  ]
},
  'networkInterfaces': [{
    'subnetwork': context.properties['control_subnet']
  }],
  'serviceAccounts': [{
    'email': context.properties['service_account_email'],
    'scopes': ['https://www.googleapis.com/auth/cloud-platform']
  }],
  'tags': {
    'items': [
      context.properties['infra_id'] + '-master',
    ]
  },
  'zone': context.properties['zones'][2]
}
]]

return {'resources': resources}

```

9.10.17. Wait for bootstrap completion and remove bootstrap resources in GCP

After you create all of the required infrastructure in Google Cloud Platform (GCP), wait for the bootstrap process to complete on the machines that you provisioned by using the Ignition config files that you generated with the installation program.

Prerequisites

- Configure a GCP account.
- Generate the Ignition config files for your cluster.

- Create and configure a VPC and associated subnets in GCP.
- Create and configure networking and load balancers in GCP.
- Create control plane and compute roles.
- Create the bootstrap machine.
- Create the control plane machines.

Procedure

1. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install wait-for bootstrap-complete --dir <installation_directory> \ 1
--log-level info 2
```

- 1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

- 2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

If the command exits without a **FATAL** warning, your production control plane has initialized.

2. Delete the bootstrap resources:

```
$ gcloud compute backend-services remove-backend ${INFRA_ID}-api-internal-backend-
service --region=${REGION} --instance-group=${INFRA_ID}-bootstrap-ig --instance-group-
zone=${ZONE_0}
```

```
$ gsutil rm gs://${INFRA_ID}-bootstrap-ignition/bootstrap.ign
```

```
$ gsutil rb gs://${INFRA_ID}-bootstrap-ignition
```

```
$ gcloud deployment-manager deployments delete ${INFRA_ID}-bootstrap
```

9.10.18. Creating additional worker machines in GCP

You can create worker machines in Google Cloud Platform (GCP) for your cluster to use by launching individual instances discretely or by automated processes outside the cluster, such as auto scaling groups. You can also take advantage of the built-in cluster scaling mechanisms and the machine API in OpenShift Container Platform.



NOTE

If you are installing a three-node cluster, skip this step. A three-node cluster consists of three control plane machines, which also act as compute machines.

In this example, you manually launch one instance by using the Deployment Manager template. Additional instances can be launched by including additional resources of type **06_worker.py** in the file.



NOTE

If you do not use the provided Deployment Manager template to create your worker machines, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- Configure a GCP account.
- Generate the Ignition config files for your cluster.
- Create and configure a VPC and associated subnets in GCP.
- Create and configure networking and load balancers in GCP.
- Create control plane and compute roles.
- Create the bootstrap machine.
- Create the control plane machines.

Procedure

1. Copy the template from the **Deployment Manager template for worker machines** section of this topic and save it as **06_worker.py** on your computer. This template describes the worker machines that your cluster requires.
2. Export the variables that the resource definition uses.

- a. Export the subnet that hosts the compute machines:

```
$ export COMPUTE_SUBNET=$(`gcloud compute networks subnets describe
${INFRA_ID}-worker-subnet --region=${REGION} --format json | jq -r .selfLink`)
```

- b. Export the email address for your service account:

```
$ export WORKER_SERVICE_ACCOUNT=$(`gcloud iam service-accounts list --filter
"email~^${INFRA_ID}-w@${PROJECT_NAME}." --format json | jq -r '[0].email`)
```

- c. Export the location of the compute machine Ignition config file:

```
$ export WORKER_IGNITION=`cat <installation_directory>/worker.ign`
```

3. Create a **06_worker.yaml** resource definition file:

```
$ cat <<EOF >06_worker.yaml
imports:
- path: 06_worker.py

resources:
- name: 'worker-0' 1
  type: 06_worker.py
  properties:
```

```

infra_id: '${INFRA_ID}' 2
zone: '${ZONE_0}' 3
compute_subnet: '${COMPUTE_SUBNET}' 4
image: '${CLUSTER_IMAGE}' 5
machine_type: 'n1-standard-4' 6
root_volume_size: '128'
service_account_email: '${WORKER_SERVICE_ACCOUNT}' 7
ignition: '${WORKER_IGNITION}' 8
- name: 'worker-1'
  type: 06_worker.py
  properties:
    infra_id: '${INFRA_ID}' 9
    zone: '${ZONE_1}' 10
    compute_subnet: '${COMPUTE_SUBNET}' 11
    image: '${CLUSTER_IMAGE}' 12
    machine_type: 'n1-standard-4' 13
    root_volume_size: '128'
    service_account_email: '${WORKER_SERVICE_ACCOUNT}' 14
    ignition: '${WORKER_IGNITION}' 15
EOF

```

- 1 **name** is the name of the worker machine, for example **worker-0**.
- 2 9 **infra_id** is the **INFRA_ID** infrastructure name from the extraction step.
- 3 10 **zone** is the zone to deploy the worker machine into, for example **us-central1-a**.
- 4 11 **compute_subnet** is the **selfLink** URL to the compute subnet.
- 5 12 **image** is the **selfLink** URL to the RHCOS image. ¹
- 6 13 **machine_type** is the machine type of the instance, for example **n1-standard-4**.
- 7 14 **service_account_email** is the email address for the worker service account that you created.
- 8 15 **ignition** is the contents of the **worker.ign** file.

4. Optional: If you want to launch additional instances, include additional resources of type **06_worker.py** in your **06_worker.yaml** resource definition file.
5. Create the deployment by using the **gcloud** CLI:

```
$ gcloud deployment-manager deployments create ${INFRA_ID}-worker --config
06_worker.yaml
```

1. To use a GCP Marketplace image, specify the offer to use:
 - OpenShift Container Platform:
<https://www.googleapis.com/compute/v1/projects/redhat-marketplace-public/global/images/redhat-coreos-ocp-413-x86-64-202305021736>

- OpenShift Platform Plus: <https://www.googleapis.com/compute/v1/projects/redhat-marketplace-public/global/images/redhat-coreos-opp-413-x86-64-202305021736>
- OpenShift Kubernetes Engine: <https://www.googleapis.com/compute/v1/projects/redhat-marketplace-public/global/images/redhat-coreos-oke-413-x86-64-202305021736>

9.10.18.1. Deployment Manager template for worker machines

You can use the following Deployment Manager template to deploy the worker machines that you need for your OpenShift Container Platform cluster:

Example 9.83. 06_worker.py Deployment Manager template


```
def GenerateConfig(context):

    resources = [{
        'name': context.properties['infra_id'] + '-' + context.env['name'],
        'type': 'compute.v1.instance',
        'properties': {
            'disks': [{
                'autoDelete': True,
                'boot': True,
                'initializeParams': {
                    'diskSizeGb': context.properties['root_volume_size'],
                    'sourceImage': context.properties['image']
                }
            }],
            'machineType': 'zones/' + context.properties['zone'] + '/machineTypes/' +
context.properties['machine_type'],
            'metadata': {
                'items': [{
                    'key': 'user-data',
                    'value': context.properties['ignition']
                }]
            },
            'networkInterfaces': [{
                'subnetwork': context.properties['compute_subnet']
            }],
            'serviceAccounts': [{
                'email': context.properties['service_account_email'],
                'scopes': ['https://www.googleapis.com/auth/cloud-platform']
            }],
            'tags': {
                'items': [
                    context.properties['infra_id'] + '-worker',
                ]
            },
            'zone': context.properties['zone']
        }
    ]

    return {'resources': resources}
```


9.10.19. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.



IMPORTANT

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**. To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.

To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.



NOTE

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your **PATH**.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

9.10.20. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

9.10.21. Approving the certificate signing requests for your machines

When you add machines to a cluster, two pending certificate signing requests (CSRs) are generated for each machine that you added. You must confirm that these CSRs are approved or, if necessary, approve them yourself. The client requests must be approved first, followed by the server requests.

Prerequisites

- You added machines to your cluster.

Procedure


1. Confirm that the cluster recognizes the machines:

```
$ oc get nodes
```

Example output

```
NAME      STATUS  ROLES  AGE  VERSION
master-0  Ready   master 63m  v1.29.4
master-1  Ready   master 63m  v1.29.4
master-2  Ready   master 64m  v1.29.4
```

The output lists all of the machines that you created.



NOTE

The preceding output might not include the compute nodes, also known as worker nodes, until some CSRs are approved.

2. Review the pending CSRs and ensure that you see the client requests with the **Pending** or **Approved** status for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

```

NAME      AGE   REQUESTOR                                     CONDITION
csr-8b2br 15m   system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper Pending
csr-8vnps 15m   system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper Pending
...

```

In this example, two machines are joining the cluster. You might see more approved CSRs in the list.

- If the CSRs were not approved, after all of the pending CSRs for the machines you added are in **Pending** status, approve the CSRs for your cluster machines:



NOTE

Because the CSRs rotate automatically, approve your CSRs within an hour of adding the machines to the cluster. If you do not approve them within an hour, the certificates will rotate, and more than two certificates will be present for each node. You must approve all of these certificates. After the client CSR is approved, the Kubelet creates a secondary CSR for the serving certificate, which requires manual approval. Then, subsequent serving certificate renewal requests are automatically approved by the **machine-approver** if the Kubelet requests a new certificate with identical parameters.



NOTE

For clusters running on platforms that are not machine API enabled, such as bare metal and other user-provisioned infrastructure, you must implement a method of automatically approving the kubelet serving certificate requests (CSRs). If a request is not approved, then the **oc exec**, **oc rsh**, and **oc logs** commands cannot succeed, because a serving certificate is required when the API server connects to the kubelet. Any operation that contacts the Kubelet endpoint requires this certificate approval to be in place. The method must watch for new CSRs, confirm that the CSR was submitted by the **node-bootstrapper** service account in the **system:node** or **system:admin** groups, and confirm the identity of the node.


- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{\n"}\n{{end}}\n{{end}}' | xargs --no-run-if-empty oc adm certificate approve
```

**NOTE**

Some Operators might not become available until some CSRs are approved.

- Now that your client requests are approved, you must review the server requests for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

```
NAME      AGE   REQUESTOR                                     CONDITION
csr-bfd72 5m26s system:node:ip-10-0-50-126.us-east-2.compute.internal
Pending
csr-c57lv 5m26s system:node:ip-10-0-95-157.us-east-2.compute.internal
Pending
...
```

- If the remaining CSRs are not approved, and are in the **Pending** status, approve the CSRs for your cluster machines:

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{"\n"}\n{{end}}{{end}}' | xargs oc adm certificate approve
```

- After all client and server CSRs have been approved, the machines have the **Ready** status. Verify this by running the following command:

```
$ oc get nodes
```

Example output

```
NAME      STATUS  ROLES  AGE  VERSION
master-0  Ready   master 73m  v1.29.4
master-1  Ready   master 73m  v1.29.4
master-2  Ready   master 74m  v1.29.4
worker-0  Ready   worker 11m  v1.29.4
worker-1  Ready   worker 11m  v1.29.4
```

**NOTE**

It can take a few minutes after approval of the server CSRs for the machines to transition to the **Ready** status.

Additional information

- For more information on CSRs, see [Certificate Signing Requests](#).

9.10.22. Optional: Adding the ingress DNS records

If you removed the DNS zone configuration when creating Kubernetes manifests and generating Ignition configs, you must manually create DNS records that point at the ingress load balancer. You can create either a wildcard `*.apps.{baseDomain}`, or specific records. You can use A, CNAME, and other records per your requirements.

Prerequisites

- Configure a GCP account.
- Remove the DNS Zone configuration when creating Kubernetes manifests and generating Ignition configs.
- Create and configure a VPC and associated subnets in GCP.
- Create and configure networking and load balancers in GCP.
- Create control plane and compute roles.
- Create the bootstrap machine.
- Create the control plane machines.
- Create the worker machines.

Procedure

1. Wait for the Ingress router to create a load balancer and populate the **EXTERNAL-IP** field:

```
$ oc -n openshift-ingress get service router-default
```

Example output

```
NAME          TYPE          CLUSTER-IP    EXTERNAL-IP    PORT(S)          AGE
router-default LoadBalancer  172.30.18.154  35.233.157.184 80:32288/TCP,443:31215/TCP 98
```

2. Add the A record to your zones:

- To use A records:
 - i. Export the variable for the router IP address:

```
$ export ROUTER_IP=`oc -n openshift-ingress get service router-default --no-headers | awk '{print $4}'`
```

- ii. Add the A record to the private zones:

```
$ if [ -f transaction.yaml ]; then rm transaction.yaml; fi
$ gcloud dns record-sets transaction start --zone ${INFRA_ID}-private-zone
```

```
$ gcloud dns record-sets transaction add ${ROUTER_IP} --name
\*.apps.${CLUSTER_NAME}.${BASE_DOMAIN}. --ttl 300 --type A --zone
${INFRA_ID}-private-zone
$ gcloud dns record-sets transaction execute --zone ${INFRA_ID}-private-zone
```

- iii. For an external cluster, also add the A record to the public zones:

```
$ if [ -f transaction.yaml ]; then rm transaction.yaml; fi
$ gcloud dns record-sets transaction start --zone ${BASE_DOMAIN_ZONE_NAME}
$ gcloud dns record-sets transaction add ${ROUTER_IP} --name
\*.apps.${CLUSTER_NAME}.${BASE_DOMAIN}. --ttl 300 --type A --zone
${BASE_DOMAIN_ZONE_NAME}
$ gcloud dns record-sets transaction execute --zone
${BASE_DOMAIN_ZONE_NAME}
```

- To add explicit domains instead of using a wildcard, create entries for each of the cluster's current routes:

```
$ oc get --all-namespaces -o jsonpath='{range .items[*]}{range .status.ingress[*]}{.host}
{"\n"}{end}{end}' routes
```

Example output

```
oauth-openshift.apps.your.cluster.domain.example.com
console-openshift-console.apps.your.cluster.domain.example.com
downloads-openshift-console.apps.your.cluster.domain.example.com
alertmanager-main-openshift-monitoring.apps.your.cluster.domain.example.com
prometheus-k8s-openshift-monitoring.apps.your.cluster.domain.example.com
```

9.10.23. Completing a GCP installation on user-provisioned infrastructure

After you start the OpenShift Container Platform installation on Google Cloud Platform (GCP) user-provisioned infrastructure, you can monitor the cluster events until the cluster is ready.

Prerequisites

- Deploy the bootstrap machine for an OpenShift Container Platform cluster on user-provisioned GCP infrastructure.
- Install the **oc** CLI and log in.

Procedure

1. Complete the cluster installation:

```
$ ./openshift-install --dir <installation_directory> wait-for install-complete 1
```

Example output

```
INFO Waiting up to 30m0s for the cluster to initialize...
```

- 1 For `<installation_directory>`, specify the path to the directory that you stored the installation files in.



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

2. Observe the running state of your cluster.

- a. Run the following command to view the current cluster version and status:

```
$ oc get clusterversion
```

Example output

```
NAME      VERSION  AVAILABLE  PROGRESSING  SINCE  STATUS
version   False    True        24m          Working towards 4.5.4: 99% complete
```

- b. Run the following command to view the Operators managed on the control plane by the Cluster Version Operator (CVO):

```
$ oc get clusteroperators
```

Example output

```
NAME                               VERSION  AVAILABLE  PROGRESSING  DEGRADED  SINCE
authentication                      4.5.4   True        False        False     7m56s
cloud-credential                     4.5.4   True        False        False     31m
cluster-autoscaler                   4.5.4   True        False        False     16m
console                              4.5.4   True        False        False     10m
csi-snapshot-controller              4.5.4   True        False        False     16m
dns                                  4.5.4   True        False        False     22m
etcd                                  4.5.4   False       False        False     25s
image-registry                       4.5.4   True        False        False     16m
ingress                              4.5.4   True        False        False     16m
insights                             4.5.4   True        False        False     17m
kube-apiserver                       4.5.4   True        False        False     19m
kube-controller-manager               4.5.4   True        False        False     20m
kube-scheduler                       4.5.4   True        False        False     20m
```


kube-storage-version-migrator	4.5.4	True	False	False	16m
machine-api	4.5.4	True	False	False	22m
machine-config	4.5.4	True	False	False	22m
marketplace	4.5.4	True	False	False	16m
monitoring	4.5.4	True	False	False	10m
network	4.5.4	True	False	False	23m
node-tuning	4.5.4	True	False	False	23m
openshift-apiserver	4.5.4	True	False	False	17m
openshift-controller-manager	4.5.4	True	False	False	15m
openshift-samples	4.5.4	True	False	False	16m
operator-lifecycle-manager	4.5.4	True	False	False	22m
operator-lifecycle-manager-catalog	4.5.4	True	False	False	22m
operator-lifecycle-manager-packageserver	4.5.4	True	False	False	18m
service-ca	4.5.4	True	False	False	23m
service-catalog-apiserver	4.5.4	True	False	False	23m
service-catalog-controller-manager	4.5.4	True	False	False	23m
storage	4.5.4	True	False	False	17m

- c. Run the following command to view your cluster pods:

```
$ oc get pods --all-namespaces
```

Example output

```

NAMESPACE                               NAME
READY STATUS   RESTARTS AGE
kube-system          etcd-member-ip-10-0-3-111.us-east-
2.compute.internal  1/1   Running  0    35m
kube-system          etcd-member-ip-10-0-3-239.us-east-
2.compute.internal  1/1   Running  0    37m
kube-system          etcd-member-ip-10-0-3-24.us-east-
2.compute.internal  1/1   Running  0    35m
openshift-apiserver-operator            openshift-apiserver-operator-6d6674f4f4-
h7t2t          1/1   Running  1    37m
openshift-apiserver          apiserver-fm48r
1/1   Running  0    30m
openshift-apiserver          apiserver-fxkvv
1/1   Running  0    29m
openshift-apiserver          apiserver-q85nm
1/1   Running  0    29m
...
openshift-service-ca-operator            openshift-service-ca-operator-66ff6dc6cd-
9r257          1/1   Running  0    37m
openshift-service-ca          apiservice-cabundle-injector-695b6bcbc-cl5hm
1/1   Running  0    35m
openshift-service-ca          configmap-cabundle-injector-8498544d7-
25qn6          1/1   Running  0    35m
openshift-service-ca          service-serving-cert-signer-6445fc9c6-wqdqn
1/1   Running  0    35m
openshift-service-catalog-apiserver-operator            openshift-service-catalog-apiserver-
operator-549f44668b-b5q2w  1/1   Running  0    32m
openshift-service-catalog-controller-manager-operator            openshift-service-catalog-
controller-manager-operator-b78cr2lnm  1/1   Running  0    31m

```

When the current cluster version is **AVAILABLE**, the installation is complete.

9.10.24. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service

9.10.25. Next steps

- [Customize your cluster](#).
- If necessary, you can [opt out of remote health reporting](#) .
- [Configure Global Access for an Ingress Controller on GCP](#) .

9.11. INSTALLING A CLUSTER INTO A SHARED VPC ON GCP USING DEPLOYMENT MANAGER TEMPLATES

In OpenShift Container Platform version 4.16, you can install a cluster into a shared Virtual Private Cloud (VPC) on Google Cloud Platform (GCP) that uses infrastructure that you provide. In this context, a cluster installed into a shared VPC is a cluster that is configured to use a VPC from a project different from where the cluster is being deployed.

A shared VPC enables an organization to connect resources from multiple projects to a common VPC network. You can communicate within the organization securely and efficiently by using internal IPs from that network. For more information about shared VPC, see [Shared VPC overview](#) in the GCP documentation.

The steps for performing a user-provided infrastructure installation into a shared VPC are outlined here. Several [Deployment Manager](#) templates are provided to assist in completing these steps or to help model your own. You are also free to create the required resources through other methods.




IMPORTANT

The steps for performing a user-provisioned infrastructure installation are provided as an example only. Installing a cluster with infrastructure you provide requires knowledge of the cloud provider and the installation process of OpenShift Container Platform. Several Deployment Manager templates are provided to assist in completing these steps or to help model your own. You are also free to create the required resources through other methods; the templates are just an example.

9.11.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.

- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- If you use a firewall and plan to use the Telemetry service, you [configured the firewall to allow the sites](#) that your cluster requires access to.
- If the cloud identity and access management (IAM) APIs are not accessible in your environment, or if you do not want to store an administrator-level credential secret in the **kube-system** namespace, you can [manually create and maintain long-term credentials](#).



NOTE

Be sure to also review this site list if you are configuring a proxy.

9.11.2. Certificate signing requests management

Because your cluster has limited access to automatic machine management when you use infrastructure that you provision, you must provide a mechanism for approving cluster certificate signing requests (CSRs) after installation. The **kube-controller-manager** only approves the kubelet client CSRs. The **machine-approver** cannot guarantee the validity of a serving certificate that is requested by using kubelet credentials because it cannot confirm that the correct machine issued the request. You must determine and implement a method of verifying the validity of the kubelet serving certificate requests and approving them.

9.11.3. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.



IMPORTANT

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

9.11.4. Configuring the GCP project that hosts your cluster


Before you can install OpenShift Container Platform, you must configure a Google Cloud Platform (GCP) project to host it.

9.11.4.1. Creating a GCP project

To install OpenShift Container Platform, you must create a project in your Google Cloud Platform (GCP) account to host the cluster.

Procedure

- Create a project to host your OpenShift Container Platform cluster. See [Creating and Managing Projects](#) in the GCP documentation.



IMPORTANT

Your GCP project must use the Premium Network Service Tier if you are using installer-provisioned infrastructure. The Standard Network Service Tier is not supported for clusters installed using the installation program. The installation program configures internal load balancing for the **api-int.<cluster_name>.<base_domain>** URL; the Premium Tier is required for internal load balancing.

9.11.4.2. Enabling API services in GCP

Your Google Cloud Platform (GCP) project requires access to several API services to complete OpenShift Container Platform installation.

Prerequisites

- You created a project to host your cluster.

Procedure

- Enable the following required API services in the project that hosts your cluster. You may also enable optional API services which are not required for installation. See [Enabling services](#) in the GCP documentation.

Table 9.30. Required API services

API service	Console service name
Compute Engine API	compute.googleapis.com
Cloud Resource Manager API	cloudresourcemanager.googleapis.com
Google DNS API	dns.googleapis.com
IAM Service Account Credentials API	iamcredentials.googleapis.com
Identity and Access Management (IAM) API	iam.googleapis.com
Service Usage API	serviceusage.googleapis.com

Table 9.31. Optional API services

API service	Console service name
Cloud Deployment Manager V2 API	deploymentmanager.googleapis.com
Google Cloud APIs	cloudapis.googleapis.com
Service Management API	servicemanagement.googleapis.com
Google Cloud Storage JSON API	storage-api.googleapis.com
Cloud Storage	storage-component.googleapis.com


9.11.4.3. GCP account limits

The OpenShift Container Platform cluster uses a number of Google Cloud Platform (GCP) components, but the default [Quotas](#) do not affect your ability to install a default OpenShift Container Platform cluster.

A default cluster, which contains three compute and three control plane machines, uses the following resources. Note that some resources are required only during the bootstrap process and are removed after the cluster deploys.

Table 9.32. GCP resources used in a default cluster

Service	Component	Location	Total resources required	Resources removed after bootstrap
Service account	IAM	Global	6	1
Firewall rules	Networking	Global	11	1
Forwarding rules	Compute	Global	2	0
Health checks	Compute	Global	2	0
Images	Compute	Global	1	0
Networks	Networking	Global	1	0
Routers	Networking	Global	1	0
Routes	Networking	Global	2	0
Subnetworks	Compute	Global	2	0
Target pools	Networking	Global	2	0

**NOTE**

If any of the quotas are insufficient during installation, the installation program displays an error that states both which quota was exceeded and the region.

Be sure to consider your actual cluster size, planned cluster growth, and any usage from other clusters that are associated with your account. The CPU, static IP addresses, and persistent disk SSD (storage) quotas are the ones that are most likely to be insufficient.

If you plan to deploy your cluster in one of the following regions, you will exceed the maximum storage quota and are likely to exceed the CPU quota limit:

- **asia-east2**
- **asia-northeast2**
- **asia-south1**
- **australia-southeast1**
- **europa-north1**
- **europa-west2**
- **europa-west3**
- **europa-west6**
- **northamerica-northeast1**
- **southamerica-east1**
- **us-west2**

You can increase resource quotas from the [GCP console](#), but you might need to file a support ticket. Be sure to plan your cluster size early so that you can allow time to resolve the support ticket before you install your OpenShift Container Platform cluster.

9.11.4.4. Creating a service account in GCP

OpenShift Container Platform requires a Google Cloud Platform (GCP) service account that provides authentication and authorization to access data in the Google APIs. If you do not have an existing IAM service account that contains the required roles in your project, you must create one.

Prerequisites

- You created a project to host your cluster.

Procedure


1. Create a service account in the project that you use to host your OpenShift Container Platform cluster. See [Creating a service account](#) in the GCP documentation.
2. Grant the service account the appropriate permissions. You can either grant the individual permissions that follow or assign the **Owner** role to it. See [Granting roles to a service account for specific resources](#).

**NOTE**

While making the service account an owner of the project is the easiest way to gain the required permissions, it means that service account has complete control over the project. You must determine if the risk that comes from offering that power is acceptable.

3. You can create the service account key in JSON format, or attach the service account to a GCP virtual machine. See [Creating service account keys](#) and [Creating and enabling service accounts for instances](#) in the GCP documentation.

You must have a service account key or a virtual machine with an attached service account to create the cluster.

**NOTE**

If you use a virtual machine with an attached service account to create your cluster, you must set **credentialsMode: Manual** in the **install-config.yaml** file before installation.

9.11.4.4.1. Required GCP roles

When you attach the **Owner** role to the service account that you create, you grant that service account all permissions, including those that are required to install OpenShift Container Platform. If your organization's security policies require a more restrictive set of permissions, you can create a service account with the following permissions. If you deploy your cluster into an existing virtual private cloud (VPC), the service account does not require certain networking permissions, which are noted in the following lists:

Required roles for the installation program

- Compute Admin
- Role Administrator
- Security Admin
- Service Account Admin
- Service Account Key Admin
- Service Account User
- Storage Admin

Required roles for creating network resources during installation

- DNS Administrator

Required roles for using the Cloud Credential Operator in passthrough mode

- Compute Load Balancer Admin

Required roles for user-provisioned GCP infrastructure

- Deployment Manager Editor

The following roles are applied to the service accounts that the control plane and compute machines use:

Table 9.33. GCP service account roles

Account	Roles
Control Plane	roles/compute.instanceAdmin
	roles/compute.networkAdmin
	roles/compute.securityAdmin
	roles/storage.admin
	roles/iam.serviceAccountUser
Compute	roles/compute.viewer
	roles/storage.admin

9.11.4.5. Supported GCP regions

You can deploy an OpenShift Container Platform cluster to the following Google Cloud Platform (GCP) regions:

- **africa-south1** (Johannesburg, South Africa)
- **asia-east1** (Changhua County, Taiwan)
- **asia-east2** (Hong Kong)
- **asia-northeast1** (Tokyo, Japan)
- **asia-northeast2** (Osaka, Japan)
- **asia-northeast3** (Seoul, South Korea)
- **asia-south1** (Mumbai, India)
- **asia-south2** (Delhi, India)
- **asia-southeast1** (Jurong West, Singapore)
- **asia-southeast2** (Jakarta, Indonesia)
- **australia-southeast1** (Sydney, Australia)
- **australia-southeast2** (Melbourne, Australia)
- **europa-central2** (Warsaw, Poland)
- **europa-north1** (Hamina, Finland)

- **europa-southwest1** (Madrid, Spain)
- **europa-west1** (St. Ghislain, Belgium)
- **europa-west2** (London, England, UK)
- **europa-west3** (Frankfurt, Germany)
- **europa-west4** (Eemshaven, Netherlands)
- **europa-west6** (Zürich, Switzerland)
- **europa-west8** (Milan, Italy)
- **europa-west9** (Paris, France)
- **europa-west12** (Turin, Italy)
- **me-central1** (Doha, Qatar, Middle East)
- **me-central2** (Dammam, Saudi Arabia, Middle East)
- **me-west1** (Tel Aviv, Israel)
- **northamerica-northeast1** (Montréal, Québec, Canada)
- **northamerica-northeast2** (Toronto, Ontario, Canada)
- **southamerica-east1** (São Paulo, Brazil)
- **southamerica-west1** (Santiago, Chile)
- **us-central1** (Council Bluffs, Iowa, USA)
- **us-east1** (Moncks Corner, South Carolina, USA)
- **us-east4** (Ashburn, Northern Virginia, USA)
- **us-east5** (Columbus, Ohio)
- **us-south1** (Dallas, Texas)
- **us-west1** (The Dalles, Oregon, USA)
- **us-west2** (Los Angeles, California, USA)
- **us-west3** (Salt Lake City, Utah, USA)
- **us-west4** (Las Vegas, Nevada, USA)



NOTE

To determine which machine type instances are available by region and zone, see the Google [documentation](#).

9.11.4.6. Installing and configuring CLI tools for GCP

To install OpenShift Container Platform on Google Cloud Platform (GCP) using user-provisioned infrastructure, you must install and configure the CLI tools for GCP.

Prerequisites

- You created a project to host your cluster.
- You created a service account and granted it the required permissions.

Procedure

1. Install the following binaries in **\$PATH**:

- **gcloud**
- **gsutil**

See [Install the latest Cloud SDK version](#) in the GCP documentation.

2. Authenticate using the **gcloud** tool with your configured service account.
See [Authorizing with a service account](#) in the GCP documentation.

9.11.5. Requirements for a cluster with user-provisioned infrastructure

For a cluster that contains user-provisioned infrastructure, you must deploy all of the required machines.

This section describes the requirements for deploying OpenShift Container Platform on user-provisioned infrastructure.

9.11.5.1. Required machines for cluster installation

The smallest OpenShift Container Platform clusters require the following hosts:

Table 9.34. Minimum required hosts

Hosts	Description
One temporary bootstrap machine	The cluster requires the bootstrap machine to deploy the OpenShift Container Platform cluster on the three control plane machines. You can remove the bootstrap machine after you install the cluster.
Three control plane machines	The control plane machines run the Kubernetes and OpenShift Container Platform services that form the control plane.
At least two compute machines, which are also known as worker machines.	The workloads requested by OpenShift Container Platform users run on the compute machines.



IMPORTANT

To maintain high availability of your cluster, use separate physical hosts for these cluster machines.

The bootstrap and control plane machines must use Red Hat Enterprise Linux CoreOS (RHCOS) as the operating system. However, the compute machines can choose between Red Hat Enterprise Linux CoreOS (RHCOS), Red Hat Enterprise Linux (RHEL) 8.6 and later.

Note that RHCOS is based on Red Hat Enterprise Linux (RHEL) 9.2 and inherits all of its hardware certifications and requirements. See [Red Hat Enterprise Linux technology capabilities and limits](#).

9.11.5.2. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 9.35. Minimum resource requirements

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Bootstrap	RHCOS	4	16 GB	100 GB	300
Control plane	RHCOS	4	16 GB	100 GB	300
Compute	RHCOS, RHEL 8.6 and later [3]	2	8 GB	100 GB	300

1. One vCPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio: $(\text{threads per core} \times \text{cores}) \times \text{sockets} = \text{vCPUs}$.
2. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes which require a 10 ms p99 fsync duration. Note that on many cloud platforms, storage size and IOPS scale together, so you might need to over-allocate storage volume to obtain sufficient performance.
3. As with all user-provisioned installations, if you choose to use RHEL compute machines in your cluster, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Use of RHEL 7 compute machines is deprecated and has been removed in OpenShift Container Platform 4.10 and later.

**NOTE**

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

9.11.5.3. Tested instance types for GCP

The following Google Cloud Platform instance types have been tested with OpenShift Container Platform.

Example 9.84. Machine series

- **A2**
- **A3**
- **C2**
- **C2D**
- **C3**
- **C3D**
- **E2**
- **M1**
- **N1**
- **N2**
- **N2D**
- **Tau T2D**

9.11.5.4. Using custom machine types

Using a custom machine type to install a OpenShift Container Platform cluster is supported.

Consider the following when using a custom machine type:

- Similar to predefined instance types, custom machine types must meet the minimum resource requirements for control plane and compute machines. For more information, see "Minimum resource requirements for cluster installation".
- The name of the custom machine type must adhere to the following syntax:
custom-<number_of_cpus>-<amount_of_memory_in_mb>

For example, **custom-6-20480**.

9.11.6. Configuring the GCP project that hosts your shared VPC network

If you use a shared Virtual Private Cloud (VPC) to host your OpenShift Container Platform cluster in Google Cloud Platform (GCP), you must configure the project that hosts it.



NOTE

If you already have a project that hosts the shared VPC network, review this section to ensure that the project meets all of the requirements to install an OpenShift Container Platform cluster.

Procedure

1. Create a project to host the shared VPC for your OpenShift Container Platform cluster. See [Creating and Managing Projects](#) in the GCP documentation.
2. Create a service account in the project that hosts your shared VPC. See [Creating a service account](#) in the GCP documentation.
3. Grant the service account the appropriate permissions. You can either grant the individual permissions that follow or assign the **Owner** role to it. See [Granting roles to a service account for specific resources](#).

**NOTE**

While making the service account an owner of the project is the easiest way to gain the required permissions, it means that service account has complete control over the project. You must determine if the risk that comes from offering that power is acceptable.

The service account for the project that hosts the shared VPC network requires the following roles:


- Compute Network User
- Compute Security Admin
- Deployment Manager Editor
- DNS Administrator
- Security Admin
- Network Management Admin

9.11.6.1. Configuring DNS for GCP

To install OpenShift Container Platform, the Google Cloud Platform (GCP) account you use must have a dedicated public hosted zone in the project that hosts the shared VPC that you install the cluster into. This zone must be authoritative for the domain. The DNS service provides cluster DNS resolution and name lookup for external connections to the cluster.

Procedure

1. Identify your domain, or subdomain, and registrar. You can transfer an existing domain and registrar or obtain a new one through GCP or another source.

**NOTE**

If you purchase a new domain, it can take time for the relevant DNS changes to propagate. For more information about purchasing domains through Google, see [Google Domains](#).

2. Create a public hosted zone for your domain or subdomain in your GCP project. See [Creating public zones](#) in the GCP documentation.
Use an appropriate root domain, such as **openshiftcorp.com**, or subdomain, such as **clusters.openshiftcorp.com**.
3. Extract the new authoritative name servers from the hosted zone records. See [Look up your Cloud DNS name servers](#) in the GCP documentation.
You typically have four name servers.
4. Update the registrar records for the name servers that your domain uses. For example, if you registered your domain to Google Domains, see the following topic in the Google Domains Help: [How to switch to custom name servers](#).
5. If you migrated your root domain to Google Cloud DNS, migrate your DNS records. See [Migrating to Cloud DNS](#) in the GCP documentation.

6. If you use a subdomain, follow your company's procedures to add its delegation records to the parent domain. This process might include a request to your company's IT department or the division that controls the root domain and DNS services for your company.

9.11.6.2. Creating a VPC in GCP

You must create a VPC in Google Cloud Platform (GCP) for your OpenShift Container Platform cluster to use. You can customize the VPC to meet your requirements. One way to create the VPC is to modify the provided Deployment Manager template.



NOTE

If you do not use the provided Deployment Manager template to create your GCP infrastructure, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- Configure a GCP account.

Procedure

1. Copy the template from the **Deployment Manager template for the VPC** section of this topic and save it as **01_vpc.py** on your computer. This template describes the VPC that your cluster requires.
2. Export the following variables required by the resource definition:

- a. Export the control plane CIDR:

```
$ export MASTER_SUBNET_CIDR='10.0.0.0/17'
```

- b. Export the compute CIDR:

```
$ export WORKER_SUBNET_CIDR='10.0.128.0/17'
```

- c. Export the region to deploy the VPC network and cluster to:

```
$ export REGION='<region>'
```

3. Export the variable for the ID of the project that hosts the shared VPC:

```
$ export HOST_PROJECT=<host_project>
```

4. Export the variable for the email of the service account that belongs to host project:

```
$ export HOST_PROJECT_ACCOUNT=<host_service_account_email>
```

5. Create a **01_vpc.yaml** resource definition file:

```
$ cat <<EOF >01_vpc.yaml
imports:
```

```

- path: 01_vpc.py

resources:
- name: cluster-vpc
  type: 01_vpc.py
  properties:
    infra_id: '<prefix>' 1
    region: '${REGION}' 2
    master_subnet_cidr: '${MASTER_SUBNET_CIDR}' 3
    worker_subnet_cidr: '${WORKER_SUBNET_CIDR}' 4
EOF

```

- 1** **infra_id** is the prefix of the network name.
- 2** **region** is the region to deploy the cluster into, for example **us-central1**.
- 3** **master_subnet_cidr** is the CIDR for the master subnet, for example **10.0.0.0/17**.
- 4** **worker_subnet_cidr** is the CIDR for the worker subnet, for example **10.0.128.0/17**.

6. Create the deployment by using the **gcloud** CLI:

```
$ gcloud deployment-manager deployments create <vpc_deployment_name> --config
01_vpc.yaml --project ${HOST_PROJECT} --account ${HOST_PROJECT_ACCOUNT} 1
```

- 1** For **<vpc_deployment_name>**, specify the name of the VPC to deploy.

7. Export the VPC variable that other components require:

- a. Export the name of the host project network:

```
$ export HOST_PROJECT_NETWORK=<vpc_network>
```

- b. Export the name of the host project control plane subnet:

```
$ export HOST_PROJECT_CONTROL_SUBNET=<control_plane_subnet>
```

- c. Export the name of the host project compute subnet:

```
$ export HOST_PROJECT_COMPUTE_SUBNET=<compute_subnet>
```

8. Set up the shared VPC. See [Setting up Shared VPC](#) in the GCP documentation.

9.11.6.2.1. Deployment Manager template for the VPC

You can use the following Deployment Manager template to deploy the VPC that you need for your OpenShift Container Platform cluster:

Example 9.85. 01_vpc.py Deployment Manager template

```
def GenerateConfig(context):
```



```

resources = [{
  'name': context.properties['infra_id'] + '-network',
  'type': 'compute.v1.network',
  'properties': {
    'region': context.properties['region'],
    'autoCreateSubnetworks': False
  }
}, {
  'name': context.properties['infra_id'] + '-master-subnet',
  'type': 'compute.v1.subnetwork',
  'properties': {
    'region': context.properties['region'],
    'network': '$(ref.' + context.properties['infra_id'] + '-network.selfLink)',
    'ipCidrRange': context.properties['master_subnet_cidr']
  }
}, {
  'name': context.properties['infra_id'] + '-worker-subnet',
  'type': 'compute.v1.subnetwork',
  'properties': {
    'region': context.properties['region'],
    'network': '$(ref.' + context.properties['infra_id'] + '-network.selfLink)',
    'ipCidrRange': context.properties['worker_subnet_cidr']
  }
}, {
  'name': context.properties['infra_id'] + '-router',
  'type': 'compute.v1.router',
  'properties': {
    'region': context.properties['region'],
    'network': '$(ref.' + context.properties['infra_id'] + '-network.selfLink)',
    'nats': [{
      'name': context.properties['infra_id'] + '-nat-master',
      'natIpAllocateOption': 'AUTO_ONLY',
      'minPortsPerVm': 7168,
      'sourceSubnetworkIpRangesToNat': 'LIST_OF_SUBNETWORKS',
      'subnetworks': [{
        'name': '$(ref.' + context.properties['infra_id'] + '-master-subnet.selfLink)',
        'sourceIpRangesToNat': ['ALL_IP_RANGES']
      }]
    }]
  }, {
    'name': context.properties['infra_id'] + '-nat-worker',
    'natIpAllocateOption': 'AUTO_ONLY',
    'minPortsPerVm': 512,
    'sourceSubnetworkIpRangesToNat': 'LIST_OF_SUBNETWORKS',
    'subnetworks': [{
      'name': '$(ref.' + context.properties['infra_id'] + '-worker-subnet.selfLink)',
      'sourceIpRangesToNat': ['ALL_IP_RANGES']
    }]
  }
}]
}
]

return {'resources': resources}

```

9.11.7. Creating the installation files for GCP

To install OpenShift Container Platform on Google Cloud Platform (GCP) using user-provisioned infrastructure, you must generate the files that the installation program needs to deploy your cluster and modify them so that the cluster creates only the machines that it will use. You generate and customize the **install-config.yaml** file, Kubernetes manifests, and Ignition config files. You also have the option to first set up a separate **var** partition during the preparation phases of installation.

9.11.7.1. Manually creating the installation configuration file

Installing the cluster requires that you manually create the installation configuration file.

Prerequisites

- You have an SSH public key on your local machine to provide to the installation program. The key will be used for SSH authentication onto your cluster nodes for debugging and disaster recovery.
- You have obtained the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create an installation directory to store your required installation assets in:

```
$ mkdir <installation_directory>
```



IMPORTANT

You must create a directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

2. Customize the sample **install-config.yaml** file template that is provided and save it in the **<installation_directory>**.



NOTE

You must name this configuration file **install-config.yaml**.

3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT


The **install-config.yaml** file is consumed during the next step of the installation process. You must back it up now.

Additional resources

- [Installation configuration parameters for GCP](#)

9.11.7.2. Enabling Shielded VMs

You can use Shielded VMs when installing your cluster. Shielded VMs have extra security features including secure boot, firmware and integrity monitoring, and rootkit detection. For more information, see Google's documentation on [Shielded VMs](#).



NOTE

Shielded VMs are currently not supported on clusters with 64-bit ARM infrastructures.

Prerequisites

- You have created an **install-config.yaml** file.

Procedure

- Use a text editor to edit the **install-config.yaml** file prior to deploying your cluster and add one of the following stanzas:
 - To use shielded VMs for only control plane machines:

```
controlPlane:
  platform:
    gcp:
      secureBoot: Enabled
```

- To use shielded VMs for only compute machines:

```
compute:
- platform:
  gcp:
    secureBoot: Enabled
```

- To use shielded VMs for all machines:

```
platform:
  gcp:
    defaultMachinePlatform:
      secureBoot: Enabled
```

9.11.7.3. Enabling Confidential VMs

You can use Confidential VMs when installing your cluster. Confidential VMs encrypt data while it is being processed. For more information, see Google's documentation on [Confidential Computing](#). You can enable Confidential VMs and Shielded VMs at the same time, although they are not dependent on each other.



NOTE

Confidential VMs are currently not supported on 64-bit ARM architectures.

Prerequisites

- You have created an **install-config.yaml** file.

Procedure

- Use a text editor to edit the **install-config.yaml** file prior to deploying your cluster and add one of the following stanzas:

- To use confidential VMs for only control plane machines:

```
controlPlane:
  platform:
    gcp:
      confidentialCompute: Enabled 1
      type: n2d-standard-8 2
      onHostMaintenance: Terminate 3
```

- Enable confidential VMs.
- Specify a machine type that supports Confidential VMs. Confidential VMs require the N2D or C2D series of machine types. For more information on supported machine types, see [Supported operating systems and machine types](#).
- Specify the behavior of the VM during a host maintenance event, such as a hardware or software update. For a machine that uses Confidential VM, this value must be set to **Terminate**, which stops the VM. Confidential VMs do not support live VM migration.

- To use confidential VMs for only compute machines:


```
compute:
  - platform:
      gcp:
        confidentialCompute: Enabled
        type: n2d-standard-8
        onHostMaintenance: Terminate
```

- To use confidential VMs for all machines:

```
platform:
  gcp:
    defaultMachinePlatform:
      confidentialCompute: Enabled
      type: n2d-standard-8
      onHostMaintenance: Terminate
```

9.11.7.4. Sample customized **install-config.yaml** file for GCP

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.



IMPORTANT

This sample YAML file is provided for reference only. You must obtain your **install-config.yaml** file by using the installation program and modify it.


```

apiVersion: v1
baseDomain: example.com 1
controlPlane: 2
  hyperthreading: Enabled 3 4
  name: master
  platform:
    gcp:
      type: n2-standard-4
      zones:
        - us-central1-a
        - us-central1-c
      tags: 5
        - control-plane-tag1
        - control-plane-tag2
  replicas: 3
compute: 6
  - hyperthreading: Enabled 7
  name: worker
  platform:
    gcp:
      type: n2-standard-4
      zones:
        - us-central1-a
        - us-central1-c
      tags: 8
        - compute-tag1
        - compute-tag2
  replicas: 0
metadata:
  name: test-cluster
networking:
  clusterNetwork:
    - cidr: 10.128.0.0/14
    hostPrefix: 23
  machineNetwork:
    - cidr: 10.0.0.0/16
  networkType: OVNKubernetes 9
  serviceNetwork:
    - 172.30.0.0/16
platform:
  gcp:
    defaultMachinePlatform:
      tags: 10
        - global-tag1
        - global-tag2
    projectID: openshift-production 11
    region: us-central1 12
pullSecret: '{"auths": ...}'
fips: false 13
sshKey: ssh-ed25519 AAAA... 14
publish: Internal 15

```

- 1 Specify the public DNS on the host project.

- 2 6 If you do not provide these parameters and values, the installation program provides the default value.
- 3 7 The **controlPlane** section is a single mapping, but the compute section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Although both sections currently define a single machine pool, it is possible that future versions of OpenShift Container Platform will support defining multiple compute pools during installation. Only one control plane pool is used.
- 4 Whether to enable or disable simultaneous multithreading, or **hypertreading**. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores. You can disable it by setting the parameter value to **Disabled**. If you disable simultaneous multithreading in some cluster machines, you must disable it in all cluster machines.



IMPORTANT

If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance. Use larger machine types, such as **n1-standard-8**, for your machines if you disable simultaneous multithreading.

- 5 8 10 Optional: A set of network tags to apply to the control plane or compute machine sets. The **platform.gcp.defaultMachinePlatform.tags** parameter applies to both control plane and compute machines. If the **compute.platform.gcp.tags** or **controlPlane.platform.gcp.tags** parameters are set, they override the **platform.gcp.defaultMachinePlatform.tags** parameter.
- 9 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 11 Specify the main project where the VM instances reside.
- 12 Specify the region that your VPC network is in.
- 13 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.




IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 14 You can optionally provide the **sshKey** value that you use to access the machines in your cluster.

**NOTE**

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- 15 How to publish the user-facing endpoints of your cluster. Set **publish** to **Internal** to deploy a private cluster, which cannot be accessed from the internet. The default value is **External**. To use a shared VPC in a cluster that uses infrastructure that you provision, you must set **publish** to **Internal**. The installation program will no longer be able to access the public DNS zone for the base domain in the host project.

9.11.7.5. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.

**NOTE**

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.


For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure


- Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.

**NOTE**

The installation program does not support the proxy **readinessEndpoints** field.

**NOTE**

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.


**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

9.11.7.6. Creating the Kubernetes manifest and Ignition config files

Because you must modify some cluster definition files and manually start the cluster machines, you must generate the Kubernetes manifest and Ignition config files that the cluster needs to configure the machines.

The installation configuration file transforms into the Kubernetes manifests. The manifests wrap into the Ignition configuration files, which are later used to configure the cluster machines.



IMPORTANT

- The Ignition config files that the OpenShift Container Platform installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

Prerequisites

- You obtained the OpenShift Container Platform installation program.
- You created the **install-config.yaml** installation configuration file.

Procedure

1. Change to the directory that contains the OpenShift Container Platform installation program and generate the Kubernetes manifests for the cluster:

```
$ ./openshift-install create manifests --dir <installation_directory> 1
```

- 1 For **<installation_directory>**, specify the installation directory that contains the **install-config.yaml** file you created.

2. Remove the Kubernetes manifest files that define the control plane machines:

```
$ rm -f <installation_directory>/openshift/99_openshift-cluster-api_master-machines-*.yaml
```

By removing these files, you prevent the cluster from automatically generating control plane machines.

3. Remove the Kubernetes manifest files that define the control plane machine set:

```
$ rm -f <installation_directory>/openshift/99_openshift-machine-api_master-control-plane-machine-set.yaml
```

4. Remove the Kubernetes manifest files that define the worker machines:

```
$ rm -f <installation_directory>/openshift/99_openshift-cluster-api_worker-machineset-*.yaml
```

Because you create and manage the worker machines yourself, you do not need to initialize these machines.

5. Check that the **mastersSchedulable** parameter in the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** Kubernetes manifest file is set to **false**. This setting prevents pods from being scheduled on the control plane machines:
 - a. Open the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** file.
 - b. Locate the **mastersSchedulable** parameter and ensure that it is set to **false**.
 - c. Save and exit the file.
6. Remove the **privateZone** sections from the **<installation_directory>/manifests/cluster-dns-02-config.yml** DNS configuration file:

```

apiVersion: config.openshift.io/v1
kind: DNS
metadata:
  creationTimestamp: null
  name: cluster
spec:
  baseDomain: example.openshift.com
  privateZone: ❶
  id: mycluster-100419-private-zone
status: {}

```

- ❶ Remove this section completely.

7. Configure the cloud provider for your VPC.
 - a. Open the **<installation_directory>/manifests/cloud-provider-config.yaml** file.
 - b. Add the **network-project-id** parameter and set its value to the ID of project that hosts the shared VPC network.
 - c. Add the **network-name** parameter and set its value to the name of the shared VPC network that hosts the OpenShift Container Platform cluster.
 - d. Replace the value of the **subnet-name** parameter with the value of the shared VPC subnet that hosts your compute machines.

The contents of the **<installation_directory>/manifests/cloud-provider-config.yaml** resemble the following example:

```

config: |+
  [global]
  project-id      = example-project
  regional       = true
  multizone      = true
  node-tags      = opensh-ptzzx-master
  node-tags      = opensh-ptzzx-worker
  node-instance-prefix = opensh-ptzzx
  external-instance-groups-prefix = opensh-ptzzx
  network-project-id = example-shared-vpc
  network-name    = example-network
  subnet-name     = example-worker-subnet

```

8. If you deploy a cluster that is not on a private network, open the `<installation_directory>/manifests/cluster-ingress-default-ingresscontroller.yaml` file and replace the value of the `scope` parameter with `External`. The contents of the file resemble the following example:

```
apiVersion: operator.openshift.io/v1
kind: IngressController
metadata:
  creationTimestamp: null
  name: default
  namespace: openshift-ingress-operator
spec:
  endpointPublishingStrategy:
    loadBalancer:
      scope: External
      type: LoadBalancerService
status:
  availableReplicas: 0
  domain: ""
  selector: ""
```

9. To create the Ignition configuration files, run the following command from the directory that contains the installation program:

```
$ ./openshift-install create ignition-configs --dir <installation_directory> ❶
```

- ❶ For `<installation_directory>`, specify the same installation directory.

Ignition config files are created for the bootstrap, control plane, and compute nodes in the installation directory. The `kubeadmin-password` and `kubeconfig` files are created in the `./<installation_directory>/auth` directory:

```
.
├── auth
│   ├── kubeadmin-password
│   └── kubeconfig
├── bootstrap.ign
├── master.ign
├── metadata.json
└── worker.ign
```

9.11.8. Exporting common variables

9.11.8.1. Extracting the infrastructure name

The Ignition config files contain a unique cluster identifier that you can use to uniquely identify your cluster in Google Cloud Platform (GCP). The infrastructure name is also used to locate the appropriate GCP resources during an OpenShift Container Platform installation. The provided Deployment Manager templates contain references to this infrastructure name, so you must extract it.

Prerequisites

- You obtained the OpenShift Container Platform installation program and the pull secret for your cluster.
- You generated the Ignition config files for your cluster.
- You installed the **jq** package.

Procedure

- To extract and view the infrastructure name from the Ignition config file metadata, run the following command:

```
$ jq -r .infraID <installation_directory>/metadata.json 1
```

- 1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

Example output

```
openshift-vw9j6 1
```

- 1 The output of this command is your cluster name and a random string.

9.11.8.2. Exporting common variables for Deployment Manager templates

You must export a common set of variables that are used with the provided Deployment Manager templates used to assist in completing a user-provided infrastructure install on Google Cloud Platform (GCP).



NOTE

Specific Deployment Manager templates can also require additional exported variables, which are detailed in their related procedures.

Prerequisites

- Obtain the OpenShift Container Platform installation program and the pull secret for your cluster.
- Generate the Ignition config files for your cluster.
- Install the **jq** package.

Procedure

1. Export the following common variables to be used by the provided Deployment Manager templates:

```
$ export BASE_DOMAIN='<base_domain>' 1  
$ export BASE_DOMAIN_ZONE_NAME='<base_domain_zone_name>' 2  
$ export NETWORK_CIDR='10.0.0.0/16'
```

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 3
$ export CLUSTER_NAME=`jq -r .clusterName <installation_directory>/metadata.json`
$ export INFRA_ID=`jq -r .infraID <installation_directory>/metadata.json`
$ export PROJECT_NAME=`jq -r .gcp.projectID <installation_directory>/metadata.json`
```

- 1** **2** Supply the values for the host project.
- 3** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

9.11.9. Networking requirements for user-provisioned infrastructure

All the Red Hat Enterprise Linux CoreOS (RHCOS) machines require networking to be configured in **initramfs** during boot to fetch their Ignition config files.

9.11.9.1. Setting the cluster node hostnames through DHCP


On Red Hat Enterprise Linux CoreOS (RHCOS) machines, the hostname is set through NetworkManager. By default, the machines obtain their hostname through DHCP. If the hostname is not provided by DHCP, set statically through kernel arguments, or another method, it is obtained through a reverse DNS lookup. Reverse DNS lookup occurs after the network has been initialized on a node and can take time to resolve. Other system services can start prior to this and detect the hostname as **localhost** or similar. You can avoid this by using DHCP to provide the hostname for each cluster node.

Additionally, setting the hostnames through DHCP can bypass any manual DNS record name configuration errors in environments that have a DNS split-horizon implementation.

9.11.9.2. Network connectivity requirements

You must configure the network connectivity between machines to allow OpenShift Container Platform cluster components to communicate. Each machine must be able to resolve the hostnames of all other machines in the cluster.

This section provides details about the ports that are required.



IMPORTANT

In connected OpenShift Container Platform environments, all nodes are required to have internet access to pull images for platform containers and provide telemetry data to Red Hat.

Table 9.36. Ports used for all-machine to all-machine communications

Protocol	Port	Description
ICMP	N/A	Network reachability tests
TCP	1936	Metrics
	9000-9999	Host level services, including the node exporter on ports 9100-9101 and the Cluster Version Operator on port 9099 .

Protocol	Port	Description
	10250-10259	The default ports that Kubernetes reserves
UDP	4789	VXLAN
	6081	Geneve
	9000-9999	Host level services, including the node exporter on ports 9100-9101 .
	500	IPsec IKE packets
	4500	IPsec NAT-T packets
	123	Network Time Protocol (NTP) on UDP port 123 If an external NTP time server is configured, you must open UDP port 123 .
TCP/UDP	30000-32767	Kubernetes node port
ESP	N/A	IPsec Encapsulating Security Payload (ESP)

Table 9.37. Ports used for all-machine to control plane communications

Protocol	Port	Description
TCP	6443	Kubernetes API

Table 9.38. Ports used for control plane machine to control plane machine communications

Protocol	Port	Description
TCP	2379-2380	etcd server and peer ports

9.11.10. Creating load balancers in GCP

You must configure load balancers in Google Cloud Platform (GCP) for your OpenShift Container Platform cluster to use. One way to create these components is to modify the provided Deployment Manager template.



NOTE

If you do not use the provided Deployment Manager template to create your GCP infrastructure, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- Configure a GCP account.
- Generate the Ignition config files for your cluster.
- Create and configure a VPC and associated subnets in GCP.

Procedure

1. Copy the template from the **Deployment Manager template for the internal load balancer** section of this topic and save it as **02_lb_int.py** on your computer. This template describes the internal load balancing objects that your cluster requires.
2. For an external cluster, also copy the template from the **Deployment Manager template for the external load balancer** section of this topic and save it as **02_lb_ext.py** on your computer. This template describes the external load balancing objects that your cluster requires.
3. Export the variables that the deployment template uses:

- a. Export the cluster network location:

```
$ export CLUSTER_NETWORK=(`gcloud compute networks describe
${HOST_PROJECT_NETWORK} --project ${HOST_PROJECT} --account
${HOST_PROJECT_ACCOUNT} --format json | jq -r .selfLink`)
```

- b. Export the control plane subnet location:

```
$ export CONTROL_SUBNET=(`gcloud compute networks subnets describe
${HOST_PROJECT_CONTROL_SUBNET} --region=${REGION} --project
${HOST_PROJECT} --account ${HOST_PROJECT_ACCOUNT} --format json | jq -r
.selfLink`)
```

- c. Export the three zones that the cluster uses:

```
$ export ZONE_0=(`gcloud compute regions describe ${REGION} --format=json | jq -r
.zones[0] | cut -d "/" -f9`)
```

```
$ export ZONE_1=(`gcloud compute regions describe ${REGION} --format=json | jq -r
.zones[1] | cut -d "/" -f9`)
```

```
$ export ZONE_2=(`gcloud compute regions describe ${REGION} --format=json | jq -r
.zones[2] | cut -d "/" -f9`)
```

4. Create a **02_infra.yaml** resource definition file:

```
$ cat <<EOF >02_infra.yaml
imports:
- path: 02_lb_ext.py
- path: 02_lb_int.py 1
resources:
- name: cluster-lb-ext 2
  type: 02_lb_ext.py
  properties:
```

```

infra_id: '${INFRA_ID}' 3
region: '${REGION}' 4
- name: cluster-lb-int
  type: 02_lb_int.py
  properties:
    cluster_network: '${CLUSTER_NETWORK}'
    control_subnet: '${CONTROL_SUBNET}' 5
    infra_id: '${INFRA_ID}'
    region: '${REGION}'
    zones: 6
      - '${ZONE_0}'
      - '${ZONE_1}'
      - '${ZONE_2}'
EOF

```

1 **2** Required only when deploying an external cluster.

3 **infra_id** is the **INFRA_ID** infrastructure name from the extraction step.

4 **region** is the region to deploy the cluster into, for example **us-central1**.

5 **control_subnet** is the URI to the control subnet.

6 **zones** are the zones to deploy the control plane instances into, like **us-east1-b**, **us-east1-c**, and **us-east1-d**.

5. Create the deployment by using the **gcloud** CLI:

```
$ gcloud deployment-manager deployments create ${INFRA_ID}-infra --config 02_infra.yaml
```

6. Export the cluster IP address:

```
$ export CLUSTER_IP=(`gcloud compute addresses describe ${INFRA_ID}-cluster-ip --
region=${REGION} --format json | jq -r .address`)
```

7. For an external cluster, also export the cluster public IP address:

```
$ export CLUSTER_PUBLIC_IP=(`gcloud compute addresses describe ${INFRA_ID}-cluster-
public-ip --region=${REGION} --format json | jq -r .address`)
```

9.11.10.1. Deployment Manager template for the external load balancer

You can use the following Deployment Manager template to deploy the external load balancer that you need for your OpenShift Container Platform cluster:

Example 9.86. 02_lb_ext.py Deployment Manager template

```

def GenerateConfig(context):

    resources = [{
        'name': context.properties['infra_id'] + '-cluster-public-ip',
        'type': 'compute.v1.address',
        'properties': {

```



```

        'region': context.properties['region']
    }
}, {
    # Refer to docs/dev/kube-apiserver-health-check.md on how to correctly setup health check
    # probe for kube-apiserver
    'name': context.properties['infra_id'] + '-api-http-health-check',
    'type': 'compute.v1.httpHealthCheck',
    'properties': {
        'port': 6080,
        'requestPath': '/readyz'
    }
}, {
    'name': context.properties['infra_id'] + '-api-target-pool',
    'type': 'compute.v1.targetPool',
    'properties': {
        'region': context.properties['region'],
        'healthChecks': ['$$(ref.' + context.properties['infra_id'] + '-api-http-health-check.selfLink)'],
        'instances': []
    }
}, {
    'name': context.properties['infra_id'] + '-api-forwarding-rule',
    'type': 'compute.v1.forwardingRule',
    'properties': {
        'region': context.properties['region'],
        'IPAddress': '$$(ref.' + context.properties['infra_id'] + '-cluster-public-ip.selfLink)',
        'target': '$$(ref.' + context.properties['infra_id'] + '-api-target-pool.selfLink)',
        'portRange': '6443'
    }
}
]]

return {'resources': resources}

```

9.11.10.2. Deployment Manager template for the internal load balancer

You can use the following Deployment Manager template to deploy the internal load balancer that you need for your OpenShift Container Platform cluster:

Example 9.87. 02_lb_int.py Deployment Manager template

```

def GenerateConfig(context):

    backends = []
    for zone in context.properties['zones']:
        backends.append({
            'group': '$$(ref.' + context.properties['infra_id'] + '-master-' + zone + '-ig' + '.selfLink)'
        })

    resources = [{
        'name': context.properties['infra_id'] + '-cluster-ip',
        'type': 'compute.v1.address',
        'properties': {
            'addressType': 'INTERNAL',
            'region': context.properties['region'],
            'subnetwork': context.properties['control_subnet']
        }
    }

```

```

    }
  }, {
    # Refer to docs/dev/kube-apiserver-health-check.md on how to correctly setup health check
    # probe for kube-apiserver
    'name': context.properties['infra_id'] + '-api-internal-health-check',
    'type': 'compute.v1.healthCheck',
    'properties': {
      'httpsHealthCheck': {
        'port': 6443,
        'requestPath': '/readyz'
      },
      'type': "HTTPS"
    }
  }, {
    'name': context.properties['infra_id'] + '-api-internal-backend-service',
    'type': 'compute.v1.regionBackendService',
    'properties': {
      'backends': backends,
      'healthChecks': ['$$(ref.' + context.properties['infra_id'] + '-api-internal-health-
check.selfLink)'],
      'loadBalancingScheme': 'INTERNAL',
      'region': context.properties['region'],
      'protocol': 'TCP',
      'timeoutSec': 120
    }
  }, {
    'name': context.properties['infra_id'] + '-api-internal-forwarding-rule',
    'type': 'compute.v1.forwardingRule',
    'properties': {
      'backendService': '$$(ref.' + context.properties['infra_id'] + '-api-internal-backend-
service.selfLink)',
      'IPAddress': '$$(ref.' + context.properties['infra_id'] + '-cluster-ip.selfLink)',
      'loadBalancingScheme': 'INTERNAL',
      'ports': ['6443','22623'],
      'region': context.properties['region'],
      'subnetwork': context.properties['control_subnet']
    }
  }
}]

for zone in context.properties['zones']:
  resources.append({
    'name': context.properties['infra_id'] + '-master-' + zone + '-ig',
    'type': 'compute.v1.instanceGroup',
    'properties': {
      'namedPorts': [
        {
          'name': 'ignition',
          'port': 22623
        }, {
          'name': 'https',
          'port': 6443
        }
      ],
      'network': context.properties['cluster_network'],
      'zone': zone
    }
  })

```

```

    })
    return {'resources': resources}

```

You will need this template in addition to the **02_lb_ext.py** template when you create an external cluster.

9.11.11. Creating a private DNS zone in GCP

You must configure a private DNS zone in Google Cloud Platform (GCP) for your OpenShift Container Platform cluster to use. One way to create this component is to modify the provided Deployment Manager template.



NOTE

If you do not use the provided Deployment Manager template to create your GCP infrastructure, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- Configure a GCP account.
- Generate the Ignition config files for your cluster.
- Create and configure a VPC and associated subnets in GCP.

Procedure

1. Copy the template from the **Deployment Manager template for the private DNS** section of this topic and save it as **02_dns.py** on your computer. This template describes the private DNS objects that your cluster requires.
2. Create a **02_dns.yaml** resource definition file:

```

$ cat <<EOF >02_dns.yaml
imports:
- path: 02_dns.py

resources:
- name: cluster-dns
  type: 02_dns.py
  properties:
    infra_id: '${INFRA_ID}' 1
    cluster_domain: '${CLUSTER_NAME}.${BASE_DOMAIN}' 2
    cluster_network: '${CLUSTER_NETWORK}' 3
EOF

```

1 **infra_id** is the **INFRA_ID** infrastructure name from the extraction step.

2 **cluster_domain** is the domain for the cluster, for example **openshift.example.com**.

3 `cluster_network` is the **selfLink** URL to the cluster network.

3. Create the deployment by using the **gcloud** CLI:

```
$ gcloud deployment-manager deployments create ${INFRA_ID}-dns --config 02_dns.yaml --project ${HOST_PROJECT} --account ${HOST_PROJECT_ACCOUNT}
```

4. The templates do not create DNS entries due to limitations of Deployment Manager, so you must create them manually:

a. Add the internal DNS entries:

```
$ if [ -f transaction.yaml ]; then rm transaction.yaml; fi
$ gcloud dns record-sets transaction start --zone ${INFRA_ID}-private-zone --project ${HOST_PROJECT} --account ${HOST_PROJECT_ACCOUNT}
$ gcloud dns record-sets transaction add ${CLUSTER_IP} --name api.${CLUSTER_NAME}.${BASE_DOMAIN}. --ttl 60 --type A --zone ${INFRA_ID}-private-zone --project ${HOST_PROJECT} --account ${HOST_PROJECT_ACCOUNT}
$ gcloud dns record-sets transaction add ${CLUSTER_IP} --name api-int.${CLUSTER_NAME}.${BASE_DOMAIN}. --ttl 60 --type A --zone ${INFRA_ID}-private-zone --project ${HOST_PROJECT} --account ${HOST_PROJECT_ACCOUNT}
$ gcloud dns record-sets transaction execute --zone ${INFRA_ID}-private-zone --project ${HOST_PROJECT} --account ${HOST_PROJECT_ACCOUNT}
```

b. For an external cluster, also add the external DNS entries:

```
$ if [ -f transaction.yaml ]; then rm transaction.yaml; fi
$ gcloud --account=${HOST_PROJECT_ACCOUNT} --project=${HOST_PROJECT} dns record-sets transaction start --zone ${BASE_DOMAIN_ZONE_NAME}
$ gcloud --account=${HOST_PROJECT_ACCOUNT} --project=${HOST_PROJECT} dns record-sets transaction add ${CLUSTER_PUBLIC_IP} --name api.${CLUSTER_NAME}.${BASE_DOMAIN}. --ttl 60 --type A --zone ${BASE_DOMAIN_ZONE_NAME}
$ gcloud --account=${HOST_PROJECT_ACCOUNT} --project=${HOST_PROJECT} dns record-sets transaction execute --zone ${BASE_DOMAIN_ZONE_NAME}
```

9.11.11.1. Deployment Manager template for the private DNS

You can use the following Deployment Manager template to deploy the private DNS that you need for your OpenShift Container Platform cluster:

Example 9.88. `02_dns.py` Deployment Manager template

```
def GenerateConfig(context):

    resources = [{
        'name': context.properties['infra_id'] + '-private-zone',
        'type': 'dns.v1.managedZone',
        'properties': {
            'description': '',
            'dnsName': context.properties['cluster_domain'] + '.',
            'visibility': 'private',
            'privateVisibilityConfig': {
```

```

        'networks': [{
            'networkUrl': context.properties['cluster_network']
        }]
    }
}
]]

return {'resources': resources}

```

9.11.12. Creating firewall rules in GCP

You must create firewall rules in Google Cloud Platform (GCP) for your OpenShift Container Platform cluster to use. One way to create these components is to modify the provided Deployment Manager template.



NOTE

If you do not use the provided Deployment Manager template to create your GCP infrastructure, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- Configure a GCP account.
- Generate the Ignition config files for your cluster.
- Create and configure a VPC and associated subnets in GCP.

Procedure

1. Copy the template from the **Deployment Manager template for firewall rules** section of this topic and save it as **03_firewall.py** on your computer. This template describes the security groups that your cluster requires.
2. Create a **03_firewall.yaml** resource definition file:

```

$ cat <<EOF >03_firewall.yaml
imports:
- path: 03_firewall.py

resources:
- name: cluster-firewall
  type: 03_firewall.py
  properties:
    allowed_external_cidr: '0.0.0.0/0' 1
    infra_id: '${INFRA_ID}' 2
    cluster_network: '${CLUSTER_NETWORK}' 3
    network_cidr: '${NETWORK_CIDR}' 4
EOF

```

- 1 **allowed_external_cidr** is the CIDR range that can access the cluster API and SSH to the bootstrap host. For an internal cluster, set this value to **\${NETWORK_CIDR}**.
- 2 **infra_id** is the **INFRA_ID** infrastructure name from the extraction step.
- 3 **cluster_network** is the **selfLink** URL to the cluster network.
- 4 **network_cidr** is the CIDR of the VPC network, for example **10.0.0.0/16**.

3. Create the deployment by using the **gcloud** CLI:

```
$ gcloud deployment-manager deployments create ${INFRA_ID}-firewall --config
03_firewall.yaml --project ${HOST_PROJECT} --account ${HOST_PROJECT_ACCOUNT}
```

9.11.12.1. Deployment Manager template for firewall rules

You can use the following Deployment Manager template to deploy the firewall rules that you need for your OpenShift Container Platform cluster:

Example 9.89. 03_firewall.py Deployment Manager template

```
def GenerateConfig(context):

    resources = [{
        'name': context.properties['infra_id'] + '-bootstrap-in-ssh',
        'type': 'compute.v1.firewall',
        'properties': {
            'network': context.properties['cluster_network'],
            'allowed': [{
                'IPProtocol': 'tcp',
                'ports': ['22']
            }],
            'sourceRanges': [context.properties['allowed_external_cidr']],
            'targetTags': [context.properties['infra_id'] + '-bootstrap']
        }
    ], {
        'name': context.properties['infra_id'] + '-api',
        'type': 'compute.v1.firewall',
        'properties': {
            'network': context.properties['cluster_network'],
            'allowed': [{
                'IPProtocol': 'tcp',
                'ports': ['6443']
            }],
            'sourceRanges': [context.properties['allowed_external_cidr']],
            'targetTags': [context.properties['infra_id'] + '-master']
        }
    ], {
        'name': context.properties['infra_id'] + '-health-checks',
        'type': 'compute.v1.firewall',
        'properties': {
            'network': context.properties['cluster_network'],
            'allowed': [{
                'IPProtocol': 'tcp',
```

```

        'ports': ['6080', '6443', '22624']
      }],
      'sourceRanges': ['35.191.0.0/16', '130.211.0.0/22', '209.85.152.0/22', '209.85.204.0/22'],
      'targetTags': [context.properties['infra_id'] + '-master']
    }
  }, {
    'name': context.properties['infra_id'] + '-etcd',
    'type': 'compute.v1.firewall',
    'properties': {
      'network': context.properties['cluster_network'],
      'allowed': [{
        'IPProtocol': 'tcp',
        'ports': ['2379-2380']
      }],
      'sourceTags': [context.properties['infra_id'] + '-master'],
      'targetTags': [context.properties['infra_id'] + '-master']
    }
  }, {
    'name': context.properties['infra_id'] + '-control-plane',
    'type': 'compute.v1.firewall',
    'properties': {
      'network': context.properties['cluster_network'],
      'allowed': [{
        'IPProtocol': 'tcp',
        'ports': ['10257']
      }],
      'sourceTags': [
        context.properties['infra_id'] + '-master',
        context.properties['infra_id'] + '-worker'
      ],
      'targetTags': [context.properties['infra_id'] + '-master']
    }
  }, {
    'name': context.properties['infra_id'] + '-internal-network',
    'type': 'compute.v1.firewall',
    'properties': {
      'network': context.properties['cluster_network'],
      'allowed': [{
        'IPProtocol': 'icmp'
      }],
      'sourceRanges': [context.properties['network_cidr']],
      'targetTags': [
        context.properties['infra_id'] + '-master',
        context.properties['infra_id'] + '-worker'
      ]
    }
  }, {

```

```

'name': context.properties['infra_id'] + '-internal-cluster',
'type': 'compute.v1.firewall',
'properties': {
  'network': context.properties['cluster_network'],
  'allowed': [{
    'IPProtocol': 'udp',
    'ports': ['4789', '6081']
  },{
    'IPProtocol': 'udp',
    'ports': ['500', '4500']
  },{
    'IPProtocol': 'esp',
  },{
    'IPProtocol': 'tcp',
    'ports': ['9000-9999']
  },{
    'IPProtocol': 'udp',
    'ports': ['9000-9999']
  },{
    'IPProtocol': 'tcp',
    'ports': ['10250']
  },{
    'IPProtocol': 'tcp',
    'ports': ['30000-32767']
  },{
    'IPProtocol': 'udp',
    'ports': ['30000-32767']
  }],
  'sourceTags': [
    context.properties['infra_id'] + '-master',
    context.properties['infra_id'] + '-worker'
  ],
  'targetTags': [
    context.properties['infra_id'] + '-master',
    context.properties['infra_id'] + '-worker'
  ]
}
}]

return {'resources': resources}

```

9.11.13. Creating IAM roles in GCP

You must create IAM roles in Google Cloud Platform (GCP) for your OpenShift Container Platform cluster to use. One way to create these components is to modify the provided Deployment Manager template.



NOTE

If you do not use the provided Deployment Manager template to create your GCP infrastructure, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- Configure a GCP account.
- Generate the Ignition config files for your cluster.
- Create and configure a VPC and associated subnets in GCP.

Procedure

1. Copy the template from the **Deployment Manager template for IAM roles** section of this topic and save it as **03_iam.py** on your computer. This template describes the IAM roles that your cluster requires.
2. Create a **03_iam.yaml** resource definition file:

```
$ cat <<EOF >03_iam.yaml
imports:
- path: 03_iam.py
resources:
- name: cluster-iam
  type: 03_iam.py
  properties:
    infra_id: '${INFRA_ID}' 1
EOF
```

- 1 **infra_id** is the **INFRA_ID** infrastructure name from the extraction step.

3. Create the deployment by using the **gcloud** CLI:

```
$ gcloud deployment-manager deployments create ${INFRA_ID}-iam --config 03_iam.yaml
```

4. Export the variable for the master service account:

```
$ export MASTER_SERVICE_ACCOUNT=$(gcloud iam service-accounts list --filter
"email~^${INFRA_ID}-m@${PROJECT_NAME}." --format json | jq -r '[0].email')
```

5. Export the variable for the worker service account:

```
$ export WORKER_SERVICE_ACCOUNT=$(gcloud iam service-accounts list --filter
"email~^${INFRA_ID}-w@${PROJECT_NAME}." --format json | jq -r '[0].email')
```

6. Assign the permissions that the installation program requires to the service accounts for the subnets that host the control plane and compute subnets:

- a. Grant the **networkViewer** role of the project that hosts your shared VPC to the master service account:

```
$ gcloud --account=${HOST_PROJECT_ACCOUNT} --project=${HOST_PROJECT}
projects add-iam-policy-binding ${HOST_PROJECT} --member
"serviceAccount:${MASTER_SERVICE_ACCOUNT}" --role
"roles/compute.networkViewer"
```

- b. Grant the **networkUser** role to the master service account for the control plane subnet:

```
$ gcloud --account=${HOST_PROJECT_ACCOUNT} --project=${HOST_PROJECT}
compute networks subnets add-iam-policy-binding
"${HOST_PROJECT_CONTROL_SUBNET}" --member
"serviceAccount:${MASTER_SERVICE_ACCOUNT}" --role "roles/compute.networkUser"
--region ${REGION}
```

- c. Grant the **networkUser** role to the worker service account for the control plane subnet:

```
$ gcloud --account=${HOST_PROJECT_ACCOUNT} --project=${HOST_PROJECT}
compute networks subnets add-iam-policy-binding
"${HOST_PROJECT_CONTROL_SUBNET}" --member
"serviceAccount:${WORKER_SERVICE_ACCOUNT}" --role
"roles/compute.networkUser" --region ${REGION}
```

- d. Grant the **networkUser** role to the master service account for the compute subnet:

```
$ gcloud --account=${HOST_PROJECT_ACCOUNT} --project=${HOST_PROJECT}
compute networks subnets add-iam-policy-binding
"${HOST_PROJECT_COMPUTE_SUBNET}" --member
"serviceAccount:${MASTER_SERVICE_ACCOUNT}" --role "roles/compute.networkUser"
--region ${REGION}
```

- e. Grant the **networkUser** role to the worker service account for the compute subnet:

```
$ gcloud --account=${HOST_PROJECT_ACCOUNT} --project=${HOST_PROJECT}
compute networks subnets add-iam-policy-binding
"${HOST_PROJECT_COMPUTE_SUBNET}" --member
"serviceAccount:${WORKER_SERVICE_ACCOUNT}" --role
"roles/compute.networkUser" --region ${REGION}
```

7. The templates do not create the policy bindings due to limitations of Deployment Manager, so you must create them manually:

```
$ gcloud projects add-iam-policy-binding ${PROJECT_NAME} --member
"serviceAccount:${MASTER_SERVICE_ACCOUNT}" --role "roles/compute.instanceAdmin"
$ gcloud projects add-iam-policy-binding ${PROJECT_NAME} --member
"serviceAccount:${MASTER_SERVICE_ACCOUNT}" --role "roles/compute.networkAdmin"
$ gcloud projects add-iam-policy-binding ${PROJECT_NAME} --member
"serviceAccount:${MASTER_SERVICE_ACCOUNT}" --role "roles/compute.securityAdmin"
$ gcloud projects add-iam-policy-binding ${PROJECT_NAME} --member
"serviceAccount:${MASTER_SERVICE_ACCOUNT}" --role "roles/iam.serviceAccountUser"
$ gcloud projects add-iam-policy-binding ${PROJECT_NAME} --member
"serviceAccount:${MASTER_SERVICE_ACCOUNT}" --role "roles/storage.admin"

$ gcloud projects add-iam-policy-binding ${PROJECT_NAME} --member
"serviceAccount:${WORKER_SERVICE_ACCOUNT}" --role "roles/compute.viewer"
$ gcloud projects add-iam-policy-binding ${PROJECT_NAME} --member
"serviceAccount:${WORKER_SERVICE_ACCOUNT}" --role "roles/storage.admin"
```

8. Create a service account key and store it locally for later use:

```
$ gcloud iam service-accounts keys create service-account-key.json --iam-
account=${MASTER_SERVICE_ACCOUNT}
```

9.11.13.1. Deployment Manager template for IAM roles

You can use the following Deployment Manager template to deploy the IAM roles that you need for your OpenShift Container Platform cluster:

Example 9.90. 03_iam.py Deployment Manager template

```
def GenerateConfig(context):

    resources = [{
        'name': context.properties['infra_id'] + '-master-node-sa',
        'type': 'iam.v1.serviceAccount',
        'properties': {
            'accountId': context.properties['infra_id'] + '-m',
            'displayName': context.properties['infra_id'] + '-master-node'
        }
    }, {
        'name': context.properties['infra_id'] + '-worker-node-sa',
        'type': 'iam.v1.serviceAccount',
        'properties': {
            'accountId': context.properties['infra_id'] + '-w',
            'displayName': context.properties['infra_id'] + '-worker-node'
        }
    }
    ]


    return {'resources': resources}
```

9.11.14. Creating the RHCOS cluster image for the GCP infrastructure

You must use a valid Red Hat Enterprise Linux CoreOS (RHCOS) image for Google Cloud Platform (GCP) for your OpenShift Container Platform nodes.

Procedure

1. Obtain the RHCOS image from the [RHCOS image mirror](#) page.



IMPORTANT

The RHCOS images might not change with every release of OpenShift Container Platform. You must download an image with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Use the image version that matches your OpenShift Container Platform version if it is available.

The file name contains the OpenShift Container Platform version number in the format **rhcos-
<version>-<arch>-gcp.<arch>.tar.gz**.

2. Create the Google storage bucket:

```
$ gsutil mb gs://<bucket_name>
```

3. Upload the RHCOS image to the Google storage bucket:

■

```
$ gsutil cp <downloaded_image_file_path>/rhcos-<version>-x86_64-gcp.x86_64.tar.gz
gs://<bucket_name>
```

- Export the uploaded RHCOS image location as a variable:


```
$ export IMAGE_SOURCE=gs://<bucket_name>/rhcos-<version>-x86_64-gcp.x86_64.tar.gz
```

- Create the cluster image:

```
$ gcloud compute images create "${INFRA_ID}-rhcos-image" \
--source-uri="${IMAGE_SOURCE}"
```

9.11.15. Creating the bootstrap machine in GCP

You must create the bootstrap machine in Google Cloud Platform (GCP) to use during OpenShift Container Platform cluster initialization. One way to create this machine is to modify the provided Deployment Manager template.



NOTE

If you do not use the provided Deployment Manager template to create your bootstrap machine, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- Configure a GCP account.
- Generate the Ignition config files for your cluster.
- Create and configure a VPC and associated subnets in GCP.
- Create and configure networking and load balancers in GCP.
- Create control plane and compute roles.
- Ensure pyOpenSSL is installed.

Procedure

- Copy the template from the **Deployment Manager template for the bootstrap machine** section of this topic and save it as **04_bootstrap.py** on your computer. This template describes the bootstrap machine that your cluster requires.
- Export the location of the Red Hat Enterprise Linux CoreOS (RHCOS) image that the installation program requires:

```
$ export CLUSTER_IMAGE=(`gcloud compute images describe ${INFRA_ID}-rhcos-image --
format json | jq -r .selfLink`)
```

- Create a bucket and upload the **bootstrap.ign** file:

```
$ gsutil mb gs://${INFRA_ID}-bootstrap-ignition
```

```
$ gsutil cp <installation_directory>/bootstrap.ign gs://${INFRA_ID}-bootstrap-ignition/
```

4. Create a signed URL for the bootstrap instance to use to access the Ignition config. Export the URL from the output as a variable:

```
$ export BOOTSTRAP_IGN=`gsutil signurl -d 1h service-account-key.json gs://${INFRA_ID}-bootstrap-ignition/bootstrap.ign | grep "^gs:" | awk '{print $5}'`
```

5. Create a **04_bootstrap.yaml** resource definition file:

```
$ cat <<EOF >04_bootstrap.yaml
imports:
- path: 04_bootstrap.py

resources:
- name: cluster-bootstrap
  type: 04_bootstrap.py
  properties:
    infra_id: '${INFRA_ID}' 1
    region: '${REGION}' 2
    zone: '${ZONE_0}' 3

    cluster_network: '${CLUSTER_NETWORK}' 4
    control_subnet: '${CONTROL_SUBNET}' 5
    image: '${CLUSTER_IMAGE}' 6
    machine_type: 'n1-standard-4' 7
    root_volume_size: '128' 8

    bootstrap_ign: '${BOOTSTRAP_IGN}' 9
EOF
```

- 1 **infra_id** is the **INFRA_ID** infrastructure name from the extraction step.
- 2 **region** is the region to deploy the cluster into, for example **us-central1**.
- 3 **zone** is the zone to deploy the bootstrap instance into, for example **us-central1-b**.
- 4 **cluster_network** is the **selfLink** URL to the cluster network.
- 5 **control_subnet** is the **selfLink** URL to the control subnet.
- 6 **image** is the **selfLink** URL to the RHCOS image.
- 7 **machine_type** is the machine type of the instance, for example **n1-standard-4**.
- 8 **root_volume_size** is the boot disk size for the bootstrap machine.
- 9 **bootstrap_ign** is the URL output when creating a signed URL.

6. Create the deployment by using the **gcloud** CLI:

```
$ gcloud deployment-manager deployments create ${INFRA_ID}-bootstrap --config
04_bootstrap.yaml
```

7. Add the bootstrap instance to the internal load balancer instance group:

```
$ gcloud compute instance-groups unmanaged add-instances ${INFRA_ID}-bootstrap-ig --
zone=${ZONE_0} --instances=${INFRA_ID}-bootstrap
```

8. Add the bootstrap instance group to the internal load balancer backend service:

```
$ gcloud compute backend-services add-backend ${INFRA_ID}-api-internal-backend-service
--region=${REGION} --instance-group=${INFRA_ID}-bootstrap-ig --instance-group-
zone=${ZONE_0}
```

9.11.15.1. Deployment Manager template for the bootstrap machine

You can use the following Deployment Manager template to deploy the bootstrap machine that you need for your OpenShift Container Platform cluster:

Example 9.91. 04_bootstrap.py Deployment Manager template

```
def GenerateConfig(context):

    resources = [{
        'name': context.properties['infra_id'] + '-bootstrap-public-ip',
        'type': 'compute.v1.address',
        'properties': {
            'region': context.properties['region']
        }
    }, {
        'name': context.properties['infra_id'] + '-bootstrap',
        'type': 'compute.v1.instance',
        'properties': {
            'disks': [{
                'autoDelete': True,
                'boot': True,
                'initializeParams': {
                    'diskSizeGb': context.properties['root_volume_size'],
                    'sourceImage': context.properties['image']
                }
            }],
            'machineType': 'zones/' + context.properties['zone'] + '/machineTypes/' +
context.properties['machine_type'],
            'metadata': {
                'items': [{
                    'key': 'user-data',
                    'value': '{"ignition":{"config":{"replace":{"source":"' + context.properties['bootstrap_ign']
+ '"},"version":"3.2.0"}}}',
                }],
            }
        },
        'networkInterfaces': [{
            'subnetwork': context.properties['control_subnet'],
            'accessConfigs': [{
                'natIP': '$(ref.' + context.properties['infra_id'] + '-bootstrap-public-ip.address)'
```

```

    ]
  },
  'tags': {
    'items': [
      context.properties['infra_id'] + '-master',
      context.properties['infra_id'] + '-bootstrap'
    ]
  },
  'zone': context.properties['zone']
}
}, {
  'name': context.properties['infra_id'] + '-bootstrap-ig',
  'type': 'compute.v1.instanceGroup',
  'properties': {
    'namedPorts': [
      {
        'name': 'ignition',
        'port': 22623
      }, {
        'name': 'https',
        'port': 6443
      }
    ],
    'network': context.properties['cluster_network'],
    'zone': context.properties['zone']
  }
}
]]

return {'resources': resources}

```

9.11.16. Creating the control plane machines in GCP

You must create the control plane machines in Google Cloud Platform (GCP) for your cluster to use. One way to create these machines is to modify the provided Deployment Manager template.



NOTE

If you do not use the provided Deployment Manager template to create your control plane machines, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- Configure a GCP account.
- Generate the Ignition config files for your cluster.
- Create and configure a VPC and associated subnets in GCP.
- Create and configure networking and load balancers in GCP.
- Create control plane and compute roles.

- Create the bootstrap machine.

Procedure

1. Copy the template from the **Deployment Manager template for control plane machines** section of this topic and save it as **05_control_plane.py** on your computer. This template describes the control plane machines that your cluster requires.
2. Export the following variable required by the resource definition:

```
$ export MASTER_IGNITION=`cat <installation_directory>/master.ign`
```

3. Create a **05_control_plane.yaml** resource definition file:

```
$ cat <<EOF >05_control_plane.yaml
imports:
- path: 05_control_plane.py

resources:
- name: cluster-control-plane
  type: 05_control_plane.py
  properties:
    infra_id: '${INFRA_ID}' 1
    zones: 2
    - '${ZONE_0}'
    - '${ZONE_1}'
    - '${ZONE_2}'

    control_subnet: '${CONTROL_SUBNET}' 3
    image: '${CLUSTER_IMAGE}' 4
    machine_type: 'n1-standard-4' 5
    root_volume_size: '128'
    service_account_email: '${MASTER_SERVICE_ACCOUNT}' 6

    ignition: '${MASTER_IGNITION}' 7
EOF
```

- 1 **infra_id** is the **INFRA_ID** infrastructure name from the extraction step.
- 2 **zones** are the zones to deploy the control plane instances into, for example **us-central1-a**, **us-central1-b**, and **us-central1-c**.
- 3 **control_subnet** is the **selfLink** URL to the control subnet.
- 4 **image** is the **selfLink** URL to the RHCOS image.
- 5 **machine_type** is the machine type of the instance, for example **n1-standard-4**.
- 6 **service_account_email** is the email address for the master service account that you created.
- 7 **ignition** is the contents of the **master.ign** file.

4. Create the deployment by using the **gcloud** CLI:


```
$ gcloud deployment-manager deployments create ${INFRA_ID}-control-plane --config
05_control_plane.yaml
```

5. The templates do not manage load balancer membership due to limitations of Deployment Manager, so you must add the control plane machines manually.

- Run the following commands to add the control plane machines to the appropriate instance groups:

```
$ gcloud compute instance-groups unmanaged add-instances ${INFRA_ID}-master-
${ZONE_0}-ig --zone=${ZONE_0} --instances=${INFRA_ID}-master-0
```

```
$ gcloud compute instance-groups unmanaged add-instances ${INFRA_ID}-master-
${ZONE_1}-ig --zone=${ZONE_1} --instances=${INFRA_ID}-master-1
```

```
$ gcloud compute instance-groups unmanaged add-instances ${INFRA_ID}-master-
${ZONE_2}-ig --zone=${ZONE_2} --instances=${INFRA_ID}-master-2
```

- For an external cluster, you must also run the following commands to add the control plane machines to the target pools:

```
$ gcloud compute target-pools add-instances ${INFRA_ID}-api-target-pool --instances-
zone="${ZONE_0}" --instances=${INFRA_ID}-master-0
```

```
$ gcloud compute target-pools add-instances ${INFRA_ID}-api-target-pool --instances-
zone="${ZONE_1}" --instances=${INFRA_ID}-master-1
```

```
$ gcloud compute target-pools add-instances ${INFRA_ID}-api-target-pool --instances-
zone="${ZONE_2}" --instances=${INFRA_ID}-master-2
```

9.11.16.1. Deployment Manager template for control plane machines

You can use the following Deployment Manager template to deploy the control plane machines that you need for your OpenShift Container Platform cluster:

Example 9.92. 05_control_plane.py Deployment Manager template

```
def GenerateConfig(context):
    resources = [{
        'name': context.properties['infra_id'] + '-master-0',
        'type': 'compute.v1.instance',
        'properties': {
            'disks': [{
                'autoDelete': True,
                'boot': True,
                'initializeParams': {
                    'diskSizeGb': context.properties['root_volume_size'],
                    'diskType': 'zones/' + context.properties['zones'][0] + '/diskTypes/pd-ssd',
                    'sourceImage': context.properties['image']
                }
            }
        ]
    }],
```

```

    'machineType': 'zones/' + context.properties['zones'][0] + '/machineTypes/' +
context.properties['machine_type'],
    'metadata': {
      'items': [{
        'key': 'user-data',
        'value': context.properties['ignition']
      }]
    },
    'networkInterfaces': [{
      'subnetwork': context.properties['control_subnet']
    }],
    'serviceAccounts': [{
      'email': context.properties['service_account_email'],
      'scopes': ['https://www.googleapis.com/auth/cloud-platform']
    }],
    'tags': {
      'items': [
        context.properties['infra_id'] + '-master',
      ]
    },
    'zone': context.properties['zones'][0]
  }
}, {
  'name': context.properties['infra_id'] + '-master-1',
  'type': 'compute.v1.instance',
  'properties': {
    'disks': [{
      'autoDelete': True,
      'boot': True,
      'initializeParams': {
        'diskSizeGb': context.properties['root_volume_size'],
        'diskType': 'zones/' + context.properties['zones'][1] + '/diskTypes/pd-ssd',
        'sourceImage': context.properties['image']
      }
    }]
  },
  'machineType': 'zones/' + context.properties['zones'][1] + '/machineTypes/' +
context.properties['machine_type'],
  'metadata': {
    'items': [{
      'key': 'user-data',
      'value': context.properties['ignition']
    }]
  },
  'networkInterfaces': [{
    'subnetwork': context.properties['control_subnet']
  }],
  'serviceAccounts': [{
    'email': context.properties['service_account_email'],
    'scopes': ['https://www.googleapis.com/auth/cloud-platform']
  }],
  'tags': {
    'items': [
      context.properties['infra_id'] + '-master',
    ]
  },
  'zone': context.properties['zones'][1]
}

```

```

    }
  }, {
    'name': context.properties['infra_id'] + '-master-2',
    'type': 'compute.v1.instance',
    'properties': {
      'disks': [{
        'autoDelete': True,
        'boot': True,
        'initializeParams': {
          'diskSizeGb': context.properties['root_volume_size'],
          'diskType': 'zones/' + context.properties['zones'][2] + '/diskTypes/pd-ssd',
          'sourceImage': context.properties['image']
        }
      }
    ],
    'machineType': 'zones/' + context.properties['zones'][2] + '/machineTypes/' +
context.properties['machine_type'],
    'metadata': {
      'items': [{
        'key': 'user-data',
        'value': context.properties['ignition']
      }
    ]
  },
  'networkInterfaces': [{
    'subnetwork': context.properties['control_subnet']
  }],
  'serviceAccounts': [{
    'email': context.properties['service_account_email'],
    'scopes': ['https://www.googleapis.com/auth/cloud-platform']
  }],
  'tags': {
    'items': [
      context.properties['infra_id'] + '-master',
    ]
  },
  'zone': context.properties['zones'][2]
}
}]

return {'resources': resources}

```

9.11.17. Wait for bootstrap completion and remove bootstrap resources in GCP

After you create all of the required infrastructure in Google Cloud Platform (GCP), wait for the bootstrap process to complete on the machines that you provisioned by using the Ignition config files that you generated with the installation program.

Prerequisites

- Configure a GCP account.
- Generate the Ignition config files for your cluster.
- Create and configure a VPC and associated subnets in GCP.

- Create and configure networking and load balancers in GCP.
- Create control plane and compute roles.
- Create the bootstrap machine.
- Create the control plane machines.

Procedure

1. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install wait-for bootstrap-complete --dir <installation_directory> \ 1
--log-level info 2
```

1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

If the command exits without a **FATAL** warning, your production control plane has initialized.

2. Delete the bootstrap resources:

```
$ gcloud compute backend-services remove-backend ${INFRA_ID}-api-internal-backend-
service --region=${REGION} --instance-group=${INFRA_ID}-bootstrap-ig --instance-group-
zone=${ZONE_0}
```

```
$ gsutil rm gs://${INFRA_ID}-bootstrap-ignition/bootstrap.ign
```

```
$ gsutil rb gs://${INFRA_ID}-bootstrap-ignition
```

```
$ gcloud deployment-manager deployments delete ${INFRA_ID}-bootstrap
```

9.11.18. Creating additional worker machines in GCP

You can create worker machines in Google Cloud Platform (GCP) for your cluster to use by launching individual instances discretely or by automated processes outside the cluster, such as auto scaling groups. You can also take advantage of the built-in cluster scaling mechanisms and the machine API in OpenShift Container Platform.

In this example, you manually launch one instance by using the Deployment Manager template. Additional instances can be launched by including additional resources of type **06_worker.py** in the file.



NOTE

If you do not use the provided Deployment Manager template to create your worker machines, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- Configure a GCP account.
- Generate the Ignition config files for your cluster.
- Create and configure a VPC and associated subnets in GCP.
- Create and configure networking and load balancers in GCP.
- Create control plane and compute roles.
- Create the bootstrap machine.
- Create the control plane machines.

Procedure

1. Copy the template from the **Deployment Manager template for worker machines** section of this topic and save it as **06_worker.py** on your computer. This template describes the worker machines that your cluster requires.
2. Export the variables that the resource definition uses.

- a. Export the subnet that hosts the compute machines:

```
$ export COMPUTE_SUBNET=$(gcloud compute networks subnets describe
${HOST_PROJECT_COMPUTE_SUBNET} --region=${REGION} --project
${HOST_PROJECT} --account ${HOST_PROJECT_ACCOUNT} --format json | jq -r
.selfLink)
```

- b. Export the email address for your service account:

```
$ export WORKER_SERVICE_ACCOUNT=$(gcloud iam service-accounts list --filter
"email~^${INFRA_ID}-w@${PROJECT_NAME}." --format json | jq -r '[0].email')
```

- c. Export the location of the compute machine Ignition config file:

```
$ export WORKER_IGNITION=`cat <installation_directory>/worker.ign`
```

3. Create a **06_worker.yaml** resource definition file:

```
$ cat <<EOF >06_worker.yaml
imports:
- path: 06_worker.py

resources:
- name: 'worker-0' 1
  type: 06_worker.py
  properties:
    infra_id: '${INFRA_ID}' 2
    zone: '${ZONE_0}' 3
    compute_subnet: '${COMPUTE_SUBNET}' 4
    image: '${CLUSTER_IMAGE}' 5
    machine_type: 'n1-standard-4' 6
    root_volume_size: '128'
```

```

    service_account_email: '${WORKER_SERVICE_ACCOUNT}' 7
    ignition: '${WORKER_IGNITION}' 8
- name: 'worker-1'
  type: 06_worker.py
  properties:
    infra_id: '${INFRA_ID}' 9
    zone: '${ZONE_1}' 10
    compute_subnet: '${COMPUTE_SUBNET}' 11
    image: '${CLUSTER_IMAGE}' 12
    machine_type: 'n1-standard-4' 13
    root_volume_size: '128'
    service_account_email: '${WORKER_SERVICE_ACCOUNT}' 14
    ignition: '${WORKER_IGNITION}' 15
EOF

```

- 1 **name** is the name of the worker machine, for example **worker-0**.
- 2 9 **infra_id** is the **INFRA_ID** infrastructure name from the extraction step.
- 3 10 **zone** is the zone to deploy the worker machine into, for example **us-central1-a**.
- 4 11 **compute_subnet** is the **selfLink** URL to the compute subnet.
- 5 12 **image** is the **selfLink** URL to the RHCOS image. ¹
- 6 13 **machine_type** is the machine type of the instance, for example **n1-standard-4**.
- 7 14 **service_account_email** is the email address for the worker service account that you created.
- 8 15 **ignition** is the contents of the **worker.ign** file.

4. Optional: If you want to launch additional instances, include additional resources of type **06_worker.py** in your **06_worker.yaml** resource definition file.
5. Create the deployment by using the **gcloud** CLI:

```

$ gcloud deployment-manager deployments create ${INFRA_ID}-worker --config
06_worker.yaml

```

1. To use a GCP Marketplace image, specify the offer to use:
 - OpenShift Container Platform: <https://www.googleapis.com/compute/v1/projects/redhat-marketplace-public/global/images/redhat-coreos-ocp-413-x86-64-202305021736>
 - OpenShift Platform Plus: <https://www.googleapis.com/compute/v1/projects/redhat-marketplace-public/global/images/redhat-coreos-opp-413-x86-64-202305021736>
 - OpenShift Kubernetes Engine: <https://www.googleapis.com/compute/v1/projects/redhat-marketplace-public/global/images/redhat-coreos-oke-413-x86-64-202305021736>

9.11.18.1. Deployment Manager template for worker machines

You can use the following Deployment Manager template to deploy the worker machines that you need for your OpenShift Container Platform cluster:

Example 9.93. 06_worker.py Deployment Manager template


```
def GenerateConfig(context):

    resources = [{
        'name': context.properties['infra_id'] + '-' + context.env['name'],
        'type': 'compute.v1.instance',
        'properties': {
            'disks': [{
                'autoDelete': True,
                'boot': True,
                'initializeParams': {
                    'diskSizeGb': context.properties['root_volume_size'],
                    'sourceImage': context.properties['image']
                }
            }],
            'machineType': 'zones/' + context.properties['zone'] + '/machineTypes/' +
context.properties['machine_type'],
            'metadata': {
                'items': [{
                    'key': 'user-data',
                    'value': context.properties['ignition']
                }]
            },
            'networkInterfaces': [{
                'subnetwork': context.properties['compute_subnet']
            }],
            'serviceAccounts': [{
                'email': context.properties['service_account_email'],
                'scopes': ['https://www.googleapis.com/auth/cloud-platform']
            }],
            'tags': {
                'items': [
                    context.properties['infra_id'] + '-worker',
                ]
            },
            'zone': context.properties['zone']
        }
    ]

    return {'resources': resources}
```

9.11.19. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.



IMPORTANT

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**.
To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```


Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.



NOTE

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your PATH.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

9.11.20. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

-

- 1 For `<installation_directory>`, specify the path to the directory that you stored the installation files in.

2. Verify you can run `oc` commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

9.11.21. Approving the certificate signing requests for your machines

When you add machines to a cluster, two pending certificate signing requests (CSRs) are generated for each machine that you added. You must confirm that these CSRs are approved or, if necessary, approve them yourself. The client requests must be approved first, followed by the server requests.

Prerequisites

- You added machines to your cluster.

Procedure

1. Confirm that the cluster recognizes the machines:

```
$ oc get nodes
```

Example output

```
NAME      STATUS   ROLES    AGE   VERSION
master-0  Ready   master   63m   v1.29.4
master-1  Ready   master   63m   v1.29.4
master-2  Ready   master   64m   v1.29.4
```

The output lists all of the machines that you created.



NOTE

The preceding output might not include the compute nodes, also known as worker nodes, until some CSRs are approved.

2. Review the pending CSRs and ensure that you see the client requests with the **Pending** or **Approved** status for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

```
NAME      AGE   REQUESTOR                                     CONDITION
csr-8b2br  15m   system:serviceaccount:openshift-machine-config-operator:node-
```

```

bootstrapper Pending
csr-8vnps 15m system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper Pending
...

```

In this example, two machines are joining the cluster. You might see more approved CSRs in the list.

- If the CSRs were not approved, after all of the pending CSRs for the machines you added are in **Pending** status, approve the CSRs for your cluster machines:



NOTE

Because the CSRs rotate automatically, approve your CSRs within an hour of adding the machines to the cluster. If you do not approve them within an hour, the certificates will rotate, and more than two certificates will be present for each node. You must approve all of these certificates. After the client CSR is approved, the Kubelet creates a secondary CSR for the serving certificate, which requires manual approval. Then, subsequent serving certificate renewal requests are automatically approved by the **machine-approver** if the Kubelet requests a new certificate with identical parameters.



NOTE

For clusters running on platforms that are not machine API enabled, such as bare metal and other user-provisioned infrastructure, you must implement a method of automatically approving the kubelet serving certificate requests (CSRs). If a request is not approved, then the **oc exec**, **oc rsh**, and **oc logs** commands cannot succeed, because a serving certificate is required when the API server connects to the kubelet. Any operation that contacts the Kubelet endpoint requires this certificate approval to be in place. The method must watch for new CSRs, confirm that the CSR was submitted by the **node-bootstrapper** service account in the **system:node** or **system:admin** groups, and confirm the identity of the node.

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{"\n"}{{end}}{{end}}' | xargs --no-run-if-empty oc adm certificate approve
```



NOTE

Some Operators might not become available until some CSRs are approved.

- Now that your client requests are approved, you must review the server requests for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

```
NAME      AGE   REQUESTOR                                     CONDITION
csr-bfd72 5m26s system:node:ip-10-0-50-126.us-east-2.compute.internal
Pending
csr-c57lv 5m26s system:node:ip-10-0-95-157.us-east-2.compute.internal
Pending
...
```

- If the remaining CSRs are not approved, and are in the **Pending** status, approve the CSRs for your cluster machines:

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{"\n"}\n{{end}}{{end}}' | xargs oc adm certificate approve
```

- After all client and server CSRs have been approved, the machines have the **Ready** status. Verify this by running the following command:

```
$ oc get nodes
```

Example output

```
NAME      STATUS   ROLES    AGE   VERSION
master-0  Ready   master   73m   v1.29.4
master-1  Ready   master   73m   v1.29.4
master-2  Ready   master   74m   v1.29.4
worker-0  Ready   worker   11m   v1.29.4
worker-1  Ready   worker   11m   v1.29.4
```



NOTE

It can take a few minutes after approval of the server CSRs for the machines to transition to the **Ready** status.

Additional information

- For more information on CSRs, see [Certificate Signing Requests](#).

9.11.22. Adding the ingress DNS records

DNS zone configuration is removed when creating Kubernetes manifests and generating Ignition

configs. You must manually create DNS records that point at the ingress load balancer. You can create either a wildcard ***.apps.{baseDomain}**. or specific records. You can use A, CNAME, and other records per your requirements.

Prerequisites

- Configure a GCP account.
- Remove the DNS Zone configuration when creating Kubernetes manifests and generating Ignition configs.
- Create and configure a VPC and associated subnets in GCP.
- Create and configure networking and load balancers in GCP.
- Create control plane and compute roles.
- Create the bootstrap machine.
- Create the control plane machines.
- Create the worker machines.

Procedure

1. Wait for the Ingress router to create a load balancer and populate the **EXTERNAL-IP** field:

```
$ oc -n openshift-ingress get service router-default
```

Example output

```
NAME          TYPE          CLUSTER-IP    EXTERNAL-IP    PORT(S)          AGE
router-default LoadBalancer  172.30.18.154 35.233.157.184 80:32288/TCP,443:31215/TCP 98
```

2. Add the A record to your zones:

- To use A records:
 - i. Export the variable for the router IP address:

```
$ export ROUTER_IP=`oc -n openshift-ingress get service router-default --no-headers | awk '{print $4}'`
```

- ii. Add the A record to the private zones:

```
$ if [ -f transaction.yaml ]; then rm transaction.yaml; fi
$ gcloud dns record-sets transaction start --zone ${INFRA_ID}-private-zone --project
${HOST_PROJECT} --account ${HOST_PROJECT_ACCOUNT}
$ gcloud dns record-sets transaction add ${ROUTER_IP} --name
\*.apps.${CLUSTER_NAME}.${BASE_DOMAIN}. --ttl 300 --type A --zone
${INFRA_ID}-private-zone --project ${HOST_PROJECT} --account
${HOST_PROJECT_ACCOUNT}
$ gcloud dns record-sets transaction execute --zone ${INFRA_ID}-private-zone --
project ${HOST_PROJECT} --account ${HOST_PROJECT_ACCOUNT}
```

- iii. For an external cluster, also add the A record to the public zones:

```
$ if [ -f transaction.yaml ]; then rm transaction.yaml; fi
$ gcloud dns record-sets transaction start --zone ${BASE_DOMAIN_ZONE_NAME} -
-project ${HOST_PROJECT} --account ${HOST_PROJECT_ACCOUNT}
$ gcloud dns record-sets transaction add ${ROUTER_IP} --name
\*.apps.${CLUSTER_NAME}.${BASE_DOMAIN}. --ttl 300 --type A --zone
${BASE_DOMAIN_ZONE_NAME} --project ${HOST_PROJECT} --account
${HOST_PROJECT_ACCOUNT}
$ gcloud dns record-sets transaction execute --zone
${BASE_DOMAIN_ZONE_NAME} --project ${HOST_PROJECT} --account
${HOST_PROJECT_ACCOUNT}
```

- To add explicit domains instead of using a wildcard, create entries for each of the cluster's current routes:

```
$ oc get --all-namespaces -o jsonpath='{range .items[*]}{range .status.ingress[*]}{.host}
{"\n"}{end}{end}' routes
```

Example output

```
oauth-openshift.apps.your.cluster.domain.example.com
console-openshift-console.apps.your.cluster.domain.example.com
downloads-openshift-console.apps.your.cluster.domain.example.com
alertmanager-main-openshift-monitoring.apps.your.cluster.domain.example.com
prometheus-k8s-openshift-monitoring.apps.your.cluster.domain.example.com
```

9.11.23. Adding ingress firewall rules

The cluster requires several firewall rules. If you do not use a shared VPC, these rules are created by the Ingress Controller via the GCP cloud provider. When you use a shared VPC, you can either create cluster-wide firewall rules for all services now or create each rule based on events, when the cluster requests access. By creating each rule when the cluster requests access, you know exactly which firewall rules are required. By creating cluster-wide firewall rules, you can apply the same rule set across multiple clusters.

If you choose to create each rule based on events, you must create firewall rules after you provision the cluster and during the life of the cluster when the console notifies you that rules are missing. Events that are similar to the following event are displayed, and you must add the firewall rules that are required:

```
$ oc get events -n openshift-ingress --field-selector="reason=LoadBalancerManualChange"
```

Example output

```
Firewall change required by security admin: `gcloud compute firewall-rules create k8s-fw-
a26e631036a3f46cba28f8df67266d55 --network example-network --description "
{"kubernetes.io/service-name":"openshift-ingress/router-default", "kubernetes.io/service-
ip":"35.237.236.234"}" --allow tcp:443,tcp:80 --source-ranges 0.0.0.0/0 --target-tags exampl-fqzq7-
master,exampl-fqzq7-worker --project example-project`
```

If you encounter issues when creating these rule-based events, you can configure the cluster-wide firewall rules while your cluster is running.

9.11.23.1. Creating cluster-wide firewall rules for a shared VPC in GCP

You can create cluster-wide firewall rules to allow the access that the OpenShift Container Platform cluster requires.



WARNING

If you do not choose to create firewall rules based on cluster events, you must create cluster-wide firewall rules.

Prerequisites

- You exported the variables that the Deployment Manager templates require to deploy your cluster.
- You created the networking and load balancing components in GCP that your cluster requires.

Procedure

1. Add a single firewall rule to allow the Google Cloud Engine health checks to access all of the services. This rule enables the ingress load balancers to determine the health status of their instances.

```
$ gcloud compute firewall-rules create --allow='tcp:30000-32767,udp:30000-32767' --
network="${CLUSTER_NETWORK}" --source-
ranges='130.211.0.0/22,35.191.0.0/16,209.85.152.0/22,209.85.204.0/22' --target-
tags="${INFRA_ID}-master,${INFRA_ID}-worker" ${INFRA_ID}-ingress-hc --
account=${HOST_PROJECT_ACCOUNT} --project=${HOST_PROJECT}
```

2. Add a single firewall rule to allow access to all cluster services:

- For an external cluster:

```
$ gcloud compute firewall-rules create --allow='tcp:80,tcp:443' --
network="${CLUSTER_NETWORK}" --source-ranges="0.0.0.0/0" --target-
tags="${INFRA_ID}-master,${INFRA_ID}-worker" ${INFRA_ID}-ingress --
account=${HOST_PROJECT_ACCOUNT} --project=${HOST_PROJECT}
```

- For a private cluster:

```
$ gcloud compute firewall-rules create --allow='tcp:80,tcp:443' --
network="${CLUSTER_NETWORK}" --source-ranges=${NETWORK_CIDR} --target-
tags="${INFRA_ID}-master,${INFRA_ID}-worker" ${INFRA_ID}-ingress --
account=${HOST_PROJECT_ACCOUNT} --project=${HOST_PROJECT}
```

Because this rule only allows traffic on TCP ports **80** and **443**, ensure that you add all the ports that your services use.

9.11.24. Completing a GCP installation on user-provisioned infrastructure

After you start the OpenShift Container Platform installation on Google Cloud Platform (GCP) user-provisioned infrastructure, you can monitor the cluster events until the cluster is ready.

Prerequisites

- Deploy the bootstrap machine for an OpenShift Container Platform cluster on user-provisioned GCP infrastructure.
- Install the **oc** CLI and log in.

Procedure

1. Complete the cluster installation:

```
$ ./openshift-install --dir <installation_directory> wait-for install-complete 1
```

Example output

```
INFO Waiting up to 30m0s for the cluster to initialize...
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

2. Observe the running state of your cluster.
 - a. Run the following command to view the current cluster version and status:

```
$ oc get clusterversion
```

Example output

```
NAME      VERSION  AVAILABLE  PROGRESSING  SINCE   STATUS
version  False    True        24m         Working towards 4.5.4: 99% complete
```


- b. Run the following command to view the Operators managed on the control plane by the Cluster Version Operator (CVO):

```
$ oc get clusteroperators
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.5.4	True	False	False	7m56s
cloud-credential	4.5.4	True	False	False	31m
cluster-autoscaler	4.5.4	True	False	False	16m
console	4.5.4	True	False	False	10m
csi-snapshot-controller	4.5.4	True	False	False	16m
dns	4.5.4	True	False	False	22m
etcd	4.5.4	False	False	False	25s
image-registry	4.5.4	True	False	False	16m
ingress	4.5.4	True	False	False	16m
insights	4.5.4	True	False	False	17m
kube-apiserver	4.5.4	True	False	False	19m
kube-controller-manager	4.5.4	True	False	False	20m
kube-scheduler	4.5.4	True	False	False	20m
kube-storage-version-migrator	4.5.4	True	False	False	16m
machine-api	4.5.4	True	False	False	22m
machine-config	4.5.4	True	False	False	22m
marketplace	4.5.4	True	False	False	16m
monitoring	4.5.4	True	False	False	10m
network	4.5.4	True	False	False	23m
node-tuning	4.5.4	True	False	False	23m
openshift-apiserver	4.5.4	True	False	False	17m
openshift-controller-manager	4.5.4	True	False	False	15m
openshift-samples	4.5.4	True	False	False	16m
operator-lifecycle-manager	4.5.4	True	False	False	22m
operator-lifecycle-manager-catalog	4.5.4	True	False	False	22m
operator-lifecycle-manager-packageserver	4.5.4	True	False	False	18m
service-ca	4.5.4	True	False	False	23m
service-catalog-apiserver	4.5.4	True	False	False	23m
service-catalog-controller-manager	4.5.4	True	False	False	23m
storage	4.5.4	True	False	False	17m

- c. Run the following command to view your cluster pods:

```
$ oc get pods --all-namespaces
```

Example output

NAMESPACE	NAME	READY	STATUS	RESTARTS	AGE
kube-system	etcd-member-ip-10-0-3-111.us-east-2.compute.internal	1/1	Running	0	35m
kube-system	etcd-member-ip-10-0-3-239.us-east-2.compute.internal	1/1	Running	0	37m
kube-system	etcd-member-ip-10-0-3-24.us-east-2.compute.internal	1/1	Running	0	35m
openshift-apiserver-operator	openshift-apiserver-operator-6d6674f4f4-				

```

h7t2t                1/1    Running  1     37m
openshift-apiserver  1/1    Running  0     30m
openshift-apiserver  1/1    Running  0     29m
openshift-apiserver  1/1    Running  0     29m
...
openshift-service-ca-operator  openshift-service-ca-operator-66ff6dc6cd-
9r257                1/1    Running  0     37m
openshift-service-ca  1/1    Running  0     35m
openshift-service-ca  25qn6                1/1    Running  0     35m
openshift-service-ca  1/1    Running  0     35m
openshift-service-catalog-apiserver-operator  openshift-service-catalog-apiserver-
operator-549f44668b-b5q2w  1/1    Running  0     32m
openshift-service-catalog-controller-manager-operator  openshift-service-catalog-
controller-manager-operator-b78cr2lnm  1/1    Running  0     31m

```

When the current cluster version is **AVAILABLE**, the installation is complete.

9.11.25. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources


- See [About remote health monitoring](#) for more information about the Telemetry service

9.11.26. Next steps

- [Customize your cluster](#).
- If necessary, you can [opt out of remote health reporting](#).

9.12. INSTALLING A CLUSTER ON GCP IN A RESTRICTED NETWORK WITH USER-PROVISIONED INFRASTRUCTURE


In OpenShift Container Platform version 4.16, you can install a cluster on Google Cloud Platform (GCP) that uses infrastructure that you provide and an internal mirror of the installation release content.



IMPORTANT

While you can install an OpenShift Container Platform cluster by using mirrored installation release content, your cluster still requires internet access to use the GCP APIs.

The steps for performing a user-provided infrastructure install are outlined here. Several [Deployment Manager](#) templates are provided to assist in completing these steps or to help model your own. You are also free to create the required resources through other methods.




IMPORTANT

The steps for performing a user-provisioned infrastructure installation are provided as an example only. Installing a cluster with infrastructure you provide requires knowledge of the cloud provider and the installation process of OpenShift Container Platform. Several Deployment Manager templates are provided to assist in completing these steps or to help model your own. You are also free to create the required resources through other methods; the templates are just an example.

9.12.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [created a registry on your mirror host](#) and obtained the **imageContentSources** data for your version of OpenShift Container Platform.



IMPORTANT

Because the installation media is on the mirror host, you can use that computer to complete all installation steps.


- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to. While you might need to grant access to more sites, you must grant access to ***.googleapis.com** and **accounts.google.com**.
- If the cloud identity and access management (IAM) APIs are not accessible in your environment, or if you do not want to store an administrator-level credential secret in the **kube-system** namespace, you can [manually create and maintain long-term credentials](#).

9.12.2. About installations in restricted networks

In OpenShift Container Platform 4.16, you can perform an installation that does not require an active connection to the internet to obtain software components. Restricted network installations can be completed using installer-provisioned infrastructure or user-provisioned infrastructure, depending on the cloud platform to which you are installing the cluster.

If you choose to perform a restricted network installation on a cloud platform, you still require access to its cloud APIs. Some cloud functions, like Amazon Web Service's Route 53 DNS and IAM services, require internet access. Depending on your network, you might require less internet access for an installation on bare metal hardware, Nutanix, or on VMware vSphere.

To complete a restricted network installation, you must create a registry that mirrors the contents of the OpenShift image registry and contains the installation media. You can create this registry on a mirror host, which can access both the internet and your closed network, or by using other methods that meet your restrictions.



IMPORTANT

Because of the complexity of the configuration for user-provisioned installations, consider completing a standard user-provisioned infrastructure installation before you attempt a restricted network installation using user-provisioned infrastructure. Completing this test installation might make it easier to isolate and troubleshoot any issues that might arise during your installation in a restricted network.

9.12.2.1. Additional limits

Clusters in restricted networks have the following additional limitations and restrictions:

- The **ClusterVersion** status includes an **Unable to retrieve available updates** error.
- By default, you cannot use the contents of the Developer Catalog because you cannot access the required image stream tags.

9.12.3. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to obtain the images that are necessary to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.

9.12.4. Configuring your GCP project

Before you can install OpenShift Container Platform, you must configure a Google Cloud Platform (GCP) project to host it.

9.12.4.1. Creating a GCP project

To install OpenShift Container Platform, you must create a project in your Google Cloud Platform (GCP) account to host the cluster.

Procedure

- Create a project to host your OpenShift Container Platform cluster. See [Creating and Managing Projects](#) in the GCP documentation.



IMPORTANT

Your GCP project must use the Premium Network Service Tier if you are using installer-provisioned infrastructure. The Standard Network Service Tier is not supported for clusters installed using the installation program. The installation program configures internal load balancing for the **api-int.<cluster_name>.<base_domain>** URL; the Premium Tier is required for internal load balancing.

9.12.4.2. Enabling API services in GCP

Your Google Cloud Platform (GCP) project requires access to several API services to complete OpenShift Container Platform installation.

Prerequisites

- You created a project to host your cluster.

Procedure

- Enable the following required API services in the project that hosts your cluster. You may also enable optional API services which are not required for installation. See [Enabling services](#) in the GCP documentation.

Table 9.39. Required API services

API service	Console service name
Compute Engine API	compute.googleapis.com
Cloud Resource Manager API	cloudresourcemanager.googleapis.com
Google DNS API	dns.googleapis.com
IAM Service Account Credentials API	iamcredentials.googleapis.com
Identity and Access Management (IAM) API	iam.googleapis.com
Service Usage API	serviceusage.googleapis.com

Table 9.40. Optional API services


API service	Console service name
Google Cloud APIs	cloudapis.googleapis.com
Service Management API	servicemanagement.googleapis.com
Google Cloud Storage JSON API	storage-api.googleapis.com
Cloud Storage	storage-component.googleapis.com

9.12.4.3. Configuring DNS for GCP

To install OpenShift Container Platform, the Google Cloud Platform (GCP) account you use must have a dedicated public hosted zone in the same project that you host the OpenShift Container Platform cluster. This zone must be authoritative for the domain. The DNS service provides cluster DNS resolution and name lookup for external connections to the cluster.

Procedure

1. Identify your domain, or subdomain, and registrar. You can transfer an existing domain and registrar or obtain a new one through GCP or another source.



NOTE

If you purchase a new domain, it can take time for the relevant DNS changes to propagate. For more information about purchasing domains through Google, see [Google Domains](#).

2. Create a public hosted zone for your domain or subdomain in your GCP project. See [Creating public zones](#) in the GCP documentation.
Use an appropriate root domain, such as **openshiftcorp.com**, or subdomain, such as **clusters.openshiftcorp.com**.
3. Extract the new authoritative name servers from the hosted zone records. See [Look up your Cloud DNS name servers](#) in the GCP documentation.
You typically have four name servers.
4. Update the registrar records for the name servers that your domain uses. For example, if you registered your domain to Google Domains, see the following topic in the Google Domains Help: [How to switch to custom name servers](#).
5. If you migrated your root domain to Google Cloud DNS, migrate your DNS records. See [Migrating to Cloud DNS](#) in the GCP documentation.
6. If you use a subdomain, follow your company's procedures to add its delegation records to the parent domain. This process might include a request to your company's IT department or the division that controls the root domain and DNS services for your company.

9.12.4.4. GCP account limits


The OpenShift Container Platform cluster uses a number of Google Cloud Platform (GCP) components, but the default [Quotas](#) do not affect your ability to install a default OpenShift Container Platform cluster.

A default cluster, which contains three compute and three control plane machines, uses the following resources. Note that some resources are required only during the bootstrap process and are removed after the cluster deploys.

Table 9.41. GCP resources used in a default cluster

Service	Component	Location	Total resources required	Resources removed after bootstrap
Service account	IAM	Global	6	1
Firewall rules	Networking	Global	11	1
Forwarding rules	Compute	Global	2	0

Service	Component	Location	Total resources required	Resources removed after bootstrap
Health checks	Compute	Global	2	0
Images	Compute	Global	1	0
Networks	Networking	Global	1	0
Routers	Networking	Global	1	0
Routes	Networking	Global	2	0
Subnetworks	Compute	Global	2	0
Target pools	Networking	Global	2	0



NOTE

If any of the quotas are insufficient during installation, the installation program displays an error that states both which quota was exceeded and the region.

Be sure to consider your actual cluster size, planned cluster growth, and any usage from other clusters that are associated with your account. The CPU, static IP addresses, and persistent disk SSD (storage) quotas are the ones that are most likely to be insufficient.

If you plan to deploy your cluster in one of the following regions, you will exceed the maximum storage quota and are likely to exceed the CPU quota limit:

- **asia-east2**
- **asia-northeast2**
- **asia-south1**
- **australia-southeast1**
- **europa-north1**
- **europa-west2**
- **europa-west3**
- **europa-west6**
- **northamerica-northeast1**
- **southamerica-east1**
- **us-west2**

You can increase resource quotas from the [GCP console](#), but you might need to file a support ticket. Be sure to plan your cluster size early so that you can allow time to resolve the support ticket before you install your OpenShift Container Platform cluster.

9.12.4.5. Creating a service account in GCP

OpenShift Container Platform requires a Google Cloud Platform (GCP) service account that provides authentication and authorization to access data in the Google APIs. If you do not have an existing IAM service account that contains the required roles in your project, you must create one.

Prerequisites

- You created a project to host your cluster.

Procedure

1. Create a service account in the project that you use to host your OpenShift Container Platform cluster. See [Creating a service account](#) in the GCP documentation.
2. Grant the service account the appropriate permissions. You can either grant the individual permissions that follow or assign the **Owner** role to it. See [Granting roles to a service account for specific resources](#).



NOTE

While making the service account an owner of the project is the easiest way to gain the required permissions, it means that service account has complete control over the project. You must determine if the risk that comes from offering that power is acceptable.

3. You can create the service account key in JSON format, or attach the service account to a GCP virtual machine. See [Creating service account keys](#) and [Creating and enabling service accounts for instances](#) in the GCP documentation.
You must have a service account key or a virtual machine with an attached service account to create the cluster.



NOTE

If you use a virtual machine with an attached service account to create your cluster, you must set **credentialsMode: Manual** in the **install-config.yaml** file before installation.

9.12.4.6. Required GCP roles

When you attach the **Owner** role to the service account that you create, you grant that service account all permissions, including those that are required to install OpenShift Container Platform. If your organization's security policies require a more restrictive set of permissions, you can create a service account with the following permissions. If you deploy your cluster into an existing virtual private cloud (VPC), the service account does not require certain networking permissions, which are noted in the following lists:

Required roles for the installation program

- Compute Admin

- Role Administrator
- Security Admin
- Service Account Admin
- Service Account Key Admin
- Service Account User
- Storage Admin

Required roles for creating network resources during installation

- DNS Administrator

Required roles for using the Cloud Credential Operator in passthrough mode

- Compute Load Balancer Admin

Required roles for user-provisioned GCP infrastructure

- Deployment Manager Editor

The following roles are applied to the service accounts that the control plane and compute machines use:

Table 9.42. GCP service account roles

Account	Roles
Control Plane	roles/compute.instanceAdmin
	roles/compute.networkAdmin
	roles/compute.securityAdmin
	roles/storage.admin
	roles/iam.serviceAccountUser
Compute	roles/compute.viewer
	roles/storage.admin

9.12.4.7. Required GCP permissions for user-provisioned infrastructure

When you attach the **Owner** role to the service account that you create, you grant that service account all permissions, including those that are required to install OpenShift Container Platform.

If your organization's security policies require a more restrictive set of permissions, you can create [custom roles](#) with the necessary permissions. The following permissions are required for the user-provisioned infrastructure for creating and deleting the OpenShift Container Platform cluster.

Example 9.94. Required permissions for creating network resources

- **compute.addresses.create**
- **compute.addresses.createInternal**
- **compute.addresses.delete**
- **compute.addresses.get**
- **compute.addresses.list**
- **compute.addresses.use**
- **compute.addresses.useInternal**
- **compute.firewalls.create**
- **compute.firewalls.delete**
- **compute.firewalls.get**
- **compute.firewalls.list**
- **compute.forwardingRules.create**
- **compute.forwardingRules.get**
- **compute.forwardingRules.list**
- **compute.forwardingRules.setLabels**
- **compute.networks.create**
- **compute.networks.get**
- **compute.networks.list**
- **compute.networks.updatePolicy**
- **compute.routers.create**
- **compute.routers.get**
- **compute.routers.list**
- **compute.routers.update**
- **compute.routes.list**
- **compute.subnetworks.create**
- **compute.subnetworks.get**

- `compute.subnetworks.list`
- `compute.subnetworks.use`
- `compute.subnetworks.useExternalIp`

Example 9.95. Required permissions for creating load balancer resources

- `compute.regionBackendServices.create`
- `compute.regionBackendServices.get`
- `compute.regionBackendServices.list`
- `compute.regionBackendServices.update`
- `compute.regionBackendServices.use`
- `compute.targetPools.addInstance`
- `compute.targetPools.create`
- `compute.targetPools.get`
- `compute.targetPools.list`
- `compute.targetPools.removeInstance`
- `compute.targetPools.use`

Example 9.96. Required permissions for creating DNS resources

- `dns.changes.create`
- `dns.changes.get`
- `dns.managedZones.create`
- `dns.managedZones.get`
- `dns.managedZones.list`
- `dns.networks.bindPrivateDNSZone`
- `dns.resourceRecordSets.create`
- `dns.resourceRecordSets.list`
- `dns.resourceRecordSets.update`

Example 9.97. Required permissions for creating Service Account resources

- `iam.serviceAccountKeys.create`

- `iam.serviceAccountKeys.delete`
- `iam.serviceAccountKeys.get`
- `iam.serviceAccountKeys.list`
- `iam.serviceAccounts.actAs`
- `iam.serviceAccounts.create`
- `iam.serviceAccounts.delete`
- `iam.serviceAccounts.get`
- `iam.serviceAccounts.list`
- `resourcemanager.projects.get`
- `resourcemanager.projects.getIamPolicy`
- `resourcemanager.projects.setIamPolicy`

Example 9.98. Required permissions for creating compute resources

- `compute.disks.create`
- `compute.disks.get`
- `compute.disks.list`
- `compute.instanceGroups.create`
- `compute.instanceGroups.delete`
- `compute.instanceGroups.get`
- `compute.instanceGroups.list`
- `compute.instanceGroups.update`
- `compute.instanceGroups.use`
- `compute.instances.create`
- `compute.instances.delete`
- `compute.instances.get`
- `compute.instances.list`
- `compute.instances.setLabels`
- `compute.instances.setMetadata`
- `compute.instances.setServiceAccount`
- `compute.instances.setTags`

- `compute.instances.use`
- `compute.machineTypes.get`
- `compute.machineTypes.list`

Example 9.99. Required for creating storage resources

- `storage.buckets.create`
- `storage.buckets.delete`
- `storage.buckets.get`
- `storage.buckets.list`
- `storage.objects.create`
- `storage.objects.delete`
- `storage.objects.get`
- `storage.objects.list`

Example 9.100. Required permissions for creating health check resources

- `compute.healthChecks.create`
- `compute.healthChecks.get`
- `compute.healthChecks.list`
- `compute.healthChecks.useReadOnly`
- `compute.httpHealthChecks.create`
- `compute.httpHealthChecks.get`
- `compute.httpHealthChecks.list`
- `compute.httpHealthChecks.useReadOnly`

Example 9.101. Required permissions to get GCP zone and region related information

- `compute.globalOperations.get`
- `compute.regionOperations.get`
- `compute.regions.list`
- `compute.zoneOperations.get`
- `compute.zones.get`

- **compute.zones.list**

Example 9.102. Required permissions for checking services and quotas

- **monitoring.timeSeries.list**
- **serviceusage.quotas.get**
- **serviceusage.services.list**

Example 9.103. Required IAM permissions for installation

- **iam.roles.get**

Example 9.104. Required Images permissions for installation

- **compute.images.create**
- **compute.images.delete**
- **compute.images.get**
- **compute.images.list**

Example 9.105. Optional permission for running gather bootstrap

- **compute.instances.getSerialPortOutput**

Example 9.106. Required permissions for deleting network resources

- **compute.addresses.delete**
- **compute.addresses.deleteInternal**
- **compute.addresses.list**
- **compute.firewalls.delete**
- **compute.firewalls.list**
- **compute.forwardingRules.delete**
- **compute.forwardingRules.list**
- **compute.networks.delete**
- **compute.networks.list**
- **compute.networks.updatePolicy**

- `compute.routers.delete`
- `compute.routers.list`
- `compute.routes.list`
- `compute.subnetworks.delete`
- `compute.subnetworks.list`

Example 9.107. Required permissions for deleting load balancer resources

- `compute.regionBackendServices.delete`
- `compute.regionBackendServices.list`
- `compute.targetPools.delete`
- `compute.targetPools.list`

Example 9.108. Required permissions for deleting DNS resources

- `dns.changes.create`
- `dns.managedZones.delete`
- `dns.managedZones.get`
- `dns.managedZones.list`
- `dns.resourceRecordSets.delete`
- `dns.resourceRecordSets.list`

Example 9.109. Required permissions for deleting Service Account resources

- `iam.serviceAccounts.delete`
- `iam.serviceAccounts.get`
- `iam.serviceAccounts.list`
- `resourcemanager.projects.getIamPolicy`
- `resourcemanager.projects.setIamPolicy`

Example 9.110. Required permissions for deleting compute resources

- `compute.disks.delete`
- `compute.disks.list`

- **compute.instanceGroups.delete**
- **compute.instanceGroups.list**
- **compute.instances.delete**
- **compute.instances.list**
- **compute.instances.stop**
- **compute.machineTypes.list**

Example 9.111. Required for deleting storage resources

- **storage.buckets.delete**
- **storage.buckets.getIamPolicy**
- **storage.buckets.list**
- **storage.objects.delete**
- **storage.objects.list**

Example 9.112. Required permissions for deleting health check resources

- **compute.healthChecks.delete**
- **compute.healthChecks.list**
- **compute.httpHealthChecks.delete**
- **compute.httpHealthChecks.list**

Example 9.113. Required Images permissions for deletion

- **compute.images.delete**
- **compute.images.list**

Example 9.114. Required permissions to get Region related information

- **compute.regions.get**

Example 9.115. Required Deployment Manager permissions

- **deploymentmanager.deployments.create**
- **deploymentmanager.deployments.delete**

- `deploymentmanager.deployments.get`
- `deploymentmanager.deployments.list`
- `deploymentmanager.manifests.get`
- `deploymentmanager.operations.get`
- `deploymentmanager.resources.list`

Additional resources


- [Optimizing storage](#)

9.12.4.8. Supported GCP regions

You can deploy an OpenShift Container Platform cluster to the following Google Cloud Platform (GCP) regions:

- **africa-south1** (Johannesburg, South Africa)
- **asia-east1** (Changhua County, Taiwan)
- **asia-east2** (Hong Kong)
- **asia-northeast1** (Tokyo, Japan)
- **asia-northeast2** (Osaka, Japan)
- **asia-northeast3** (Seoul, South Korea)
- **asia-south1** (Mumbai, India)
- **asia-south2** (Delhi, India)
- **asia-southeast1** (Jurong West, Singapore)
- **asia-southeast2** (Jakarta, Indonesia)
- **australia-southeast1** (Sydney, Australia)
- **australia-southeast2** (Melbourne, Australia)
- **europa-central2** (Warsaw, Poland)
- **europa-north1** (Hamina, Finland)
- **europa-southwest1** (Madrid, Spain)
- **europa-west1** (St. Ghislain, Belgium)
- **europa-west2** (London, England, UK)
- **europa-west3** (Frankfurt, Germany)
- **europa-west4** (Eemshaven, Netherlands)

- **europa-west6** (Zürich, Switzerland)
- **europa-west8** (Milan, Italy)
- **europa-west9** (Paris, France)
- **europa-west12** (Turin, Italy)
- **me-central1** (Doha, Qatar, Middle East)
- **me-central2** (Dammam, Saudi Arabia, Middle East)
- **me-west1** (Tel Aviv, Israel)
- **northamerica-northeast1** (Montréal, Québec, Canada)
- **northamerica-northeast2** (Toronto, Ontario, Canada)
- **southamerica-east1** (São Paulo, Brazil)
- **southamerica-west1** (Santiago, Chile)
- **us-central1** (Council Bluffs, Iowa, USA)
- **us-east1** (Moncks Corner, South Carolina, USA)
- **us-east4** (Ashburn, Northern Virginia, USA)
- **us-east5** (Columbus, Ohio)
- **us-south1** (Dallas, Texas)
- **us-west1** (The Dalles, Oregon, USA)
- **us-west2** (Los Angeles, California, USA)
- **us-west3** (Salt Lake City, Utah, USA)
- **us-west4** (Las Vegas, Nevada, USA)

**NOTE**

To determine which machine type instances are available by region and zone, see the Google [documentation](#).

9.12.4.9. Installing and configuring CLI tools for GCP

To install OpenShift Container Platform on Google Cloud Platform (GCP) using user-provisioned infrastructure, you must install and configure the CLI tools for GCP.

Prerequisites

- You created a project to host your cluster.
- You created a service account and granted it the required permissions.

Procedure

1. Install the following binaries in **\$PATH**:

- **gcloud**
- **gsutil**

See [Install the latest Cloud SDK version](#) in the GCP documentation.

2. Authenticate using the **gcloud** tool with your configured service account.

See [Authorizing with a service account](#) in the GCP documentation.

9.12.5. Requirements for a cluster with user-provisioned infrastructure

For a cluster that contains user-provisioned infrastructure, you must deploy all of the required machines.

This section describes the requirements for deploying OpenShift Container Platform on user-provisioned infrastructure.

9.12.5.1. Required machines for cluster installation

The smallest OpenShift Container Platform clusters require the following hosts:

Table 9.43. Minimum required hosts

Hosts	Description
One temporary bootstrap machine	The cluster requires the bootstrap machine to deploy the OpenShift Container Platform cluster on the three control plane machines. You can remove the bootstrap machine after you install the cluster.
Three control plane machines	The control plane machines run the Kubernetes and OpenShift Container Platform services that form the control plane.
At least two compute machines, which are also known as worker machines.	The workloads requested by OpenShift Container Platform users run on the compute machines.



IMPORTANT

To maintain high availability of your cluster, use separate physical hosts for these cluster machines.

The bootstrap and control plane machines must use Red Hat Enterprise Linux CoreOS (RHCOS) as the operating system. However, the compute machines can choose between Red Hat Enterprise Linux CoreOS (RHCOS), Red Hat Enterprise Linux (RHEL) 8.6 and later.

Note that RHCOS is based on Red Hat Enterprise Linux (RHEL) 9.2 and inherits all of its hardware certifications and requirements. See [Red Hat Enterprise Linux technology capabilities and limits](#) .

9.12.5.2. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 9.44. Minimum resource requirements

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Bootstrap	RHCOS	4	16 GB	100 GB	300
Control plane	RHCOS	4	16 GB	100 GB	300
Compute	RHCOS, RHEL 8.6 and later [3]	2	8 GB	100 GB	300

1. One vCPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio: (threads per core × cores) × sockets = vCPUs.
2. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes which require a 10 ms p99 fsync duration. Note that on many cloud platforms, storage size and IOPS scale together, so you might need to over-allocate storage volume to obtain sufficient performance.
3. As with all user-provisioned installations, if you choose to use RHEL compute machines in your cluster, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Use of RHEL 7 compute machines is deprecated and has been removed in OpenShift Container Platform 4.10 and later.

NOTE

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

9.12.5.3. Tested instance types for GCP

The following Google Cloud Platform instance types have been tested with OpenShift Container Platform.

Example 9.116. Machine series

- **A2**
- **A3**
- **C2**
- **C2D**
- **C3**
- **C3D**
- **E2**
- **M1**
- **N1**
- **N2**
- **N2D**
- **Tau T2D**

9.12.5.4. Using custom machine types

Using a custom machine type to install a OpenShift Container Platform cluster is supported.

Consider the following when using a custom machine type:

- Similar to predefined instance types, custom machine types must meet the minimum resource requirements for control plane and compute machines. For more information, see "Minimum resource requirements for cluster installation".
- The name of the custom machine type must adhere to the following syntax:
custom-<number_of_cpus>-<amount_of_memory_in_mb>

For example, **custom-6-20480**.

9.12.6. Creating the installation files for GCP

To install OpenShift Container Platform on Google Cloud Platform (GCP) using user-provisioned infrastructure, you must generate the files that the installation program needs to deploy your cluster and modify them so that the cluster creates only the machines that it will use. You generate and customize the **install-config.yaml** file, Kubernetes manifests, and Ignition config files. You also have the option to first set up a separate **var** partition during the preparation phases of installation.

9.12.6.1. Optional: Creating a separate /var partition

It is recommended that disk partitioning for OpenShift Container Platform be left to the installer. However, there are cases where you might want to create separate partitions in a part of the filesystem that you expect to grow.

OpenShift Container Platform supports the addition of a single partition to attach storage to either the **/var** partition or a subdirectory of **/var**. For example:

- **/var/lib/containers**: Holds container-related content that can grow as more images and containers are added to a system.
- **/var/lib/etcd**: Holds data that you might want to keep separate for purposes such as performance optimization of etcd storage.
- **/var**: Holds data that you might want to keep separate for purposes such as auditing.

Storing the contents of a **/var** directory separately makes it easier to grow storage for those areas as needed and reinstall OpenShift Container Platform at a later date and keep that data intact. With this method, you will not have to pull all your containers again, nor will you have to copy massive log files when you update systems.

Because **/var** must be in place before a fresh installation of Red Hat Enterprise Linux CoreOS (RHCOS), the following procedure sets up the separate **/var** partition by creating a machine config manifest that is inserted during the **openshift-install** preparation phases of an OpenShift Container Platform installation.



IMPORTANT

If you follow the steps to create a separate **/var** partition in this procedure, it is not necessary to create the Kubernetes manifest and Ignition config files again as described later in this section.

Procedure

1. Create a directory to hold the OpenShift Container Platform installation files:

```
$ mkdir $HOME/clusterconfig
```

2. Run **openshift-install** to create a set of files in the **manifest** and **openshift** subdirectories. Answer the system questions as you are prompted:

```
$ openshift-install create manifests --dir $HOME/clusterconfig
```

Example output

```
? SSH Public Key ...  
INFO Credentials loaded from the "myprofile" profile in file "/home/myuser/.aws/credentials"  
INFO Consuming Install Config from target directory  
INFO Manifests created in: $HOME/clusterconfig/manifests and  
$HOME/clusterconfig/openshift
```

3. Optional: Confirm that the installation program created manifests in the **clusterconfig/openshift** directory:

```
$ ls $HOME/clusterconfig/openshift/
```


Example output

```
99_kubeadmin-password-secret.yaml
99_openshift-cluster-api_master-machines-0.yaml
99_openshift-cluster-api_master-machines-1.yaml
99_openshift-cluster-api_master-machines-2.yaml
...
```

4. Create a Butane config that configures the additional partition. For example, name the file **\$HOME/clusterconfig/98-var-partition.bu**, change the disk device name to the name of the storage device on the **worker** systems, and set the storage size as appropriate. This example places the **/var** directory on a separate partition:

```
variant: openshift
version: 4.16.0
metadata:
  labels:
    machineconfiguration.openshift.io/role: worker
  name: 98-var-partition
storage:
  disks:
    - device: /dev/disk/by-id/<device_name> ❶
      partitions:
        - label: var
          start_mib: <partition_start_offset> ❷
          size_mib: <partition_size> ❸
          number: 5
      filesystems:
        - device: /dev/disk/by-partlabel/var
          path: /var
          format: xfs
          mount_options: [defaults, prjquota] ❹
          with_mount_unit: true
```

- ❶ The storage device name of the disk that you want to partition.
- ❷ When adding a data partition to the boot disk, a minimum value of 25000 MiB (Mebibytes) is recommended. The root file system is automatically resized to fill all available space up to the specified offset. If no value is specified, or if the specified value is smaller than the recommended minimum, the resulting root file system will be too small, and future reinstalls of RHCOS might overwrite the beginning of the data partition.
- ❸ The size of the data partition in mebibytes.
- ❹ The **prjquota** mount option must be enabled for filesystems used for container storage.



NOTE

When creating a separate **/var** partition, you cannot use different instance types for worker nodes, if the different instance types do not have the same device name.

5. Create a manifest from the Butane config and save it to the **clusterconfig/openshift** directory. For example, run the following command:

```
$ butane $HOME/clusterconfig/98-var-partition.bu -o $HOME/clusterconfig/openshift/98-var-partition.yaml
```

6. Run **openshift-install** again to create Ignition configs from a set of files in the **manifest** and **openshift** subdirectories:

```
$ openshift-install create ignition-configs --dir $HOME/clusterconfig  
$ ls $HOME/clusterconfig/  
auth bootstrap.ign master.ign metadata.json worker.ign
```

Now you can use the Ignition config files as input to the installation procedures to install Red Hat Enterprise Linux CoreOS (RHCOS) systems.

9.12.6.2. Creating the installation configuration file

You can customize the OpenShift Container Platform cluster you install on Google Cloud Platform (GCP).

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster. For a restricted network installation, these files are on your mirror host.
- You have the **imageContentSources** values that were generated during mirror registry creation.
- You have obtained the contents of the certificate for your mirror registry.

Procedure

1. Create the **install-config.yaml** file.
 - a. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install create install-config --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the directory name to store the files that the installation program creates.

When specifying the directory:

- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
- Use an empty directory. Some installation assets, such as bootstrap X.509 certificates, have short expiration intervals, therefore you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

- b. At the prompts, provide the configuration details for your cloud:
 - i. Optional: Select an SSH key to use to access your cluster machines.

**NOTE**

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- ii. Select **gcp** as the platform to target.
 - iii. If you have not configured the service account key for your GCP account on your computer, you must obtain it from GCP and paste the contents of the file or enter the absolute path to the file.
 - iv. Select the project ID to provision the cluster in. The default value is specified by the service account that you configured.
 - v. Select the region to deploy the cluster to.
 - vi. Select the base domain to deploy the cluster to. The base domain corresponds to the public DNS zone that you created for your cluster.
 - vii. Enter a descriptive name for your cluster.
2. Edit the **install-config.yaml** file to give the additional information that is required for an installation in a restricted network.
 - a. Update the **pullSecret** value to contain the authentication information for your registry:

```
pullSecret: '{"auths":{"<mirror_host_name>:5000":{"auth": "<credentials>","email":
"you@example.com"}}}'
```

For **<mirror_host_name>**, specify the registry domain name that you specified in the certificate for your mirror registry, and for **<credentials>**, specify the base64-encoded user name and password for your mirror registry.

- b. Add the **additionalTrustBundle** parameter and value.

```
additionalTrustBundle: |
  -----BEGIN CERTIFICATE-----
  ///////////////////////////////////////////////////////////////////////////////////////////////////////////////////
  -----END CERTIFICATE-----
```

The value must be the contents of the certificate file that you used for your mirror registry. The certificate file can be an existing, trusted certificate authority, or the self-signed certificate that you generated for the mirror registry.

- c. Define the network and subnets for the VPC to install the cluster in under the parent **platform.gcp** field:

```
network: <existing_vpc>
controlPlaneSubnet: <control_plane_subnet>
computeSubnet: <compute_subnet>
```

-

For **platform.gcp.network**, specify the name for the existing Google VPC. For **platform.gcp.controlPlaneSubnet** and **platform.gcp.computeSubnet**, specify the existing subnets to deploy the control plane machines and compute machines, respectively.

d. Add the image content resources, which resemble the following YAML excerpt:

```
imageContentSources:
- mirrors:
  - <mirror_host_name>:5000/<repo_name>/release
    source: quay.io/openshift-release-dev/ocp-release
- mirrors:
  - <mirror_host_name>:5000/<repo_name>/release
    source: registry.redhat.io/ocp/release
```

For these values, use the **imageContentSources** that you recorded during mirror registry creation.

e. Optional: Set the publishing strategy to **Internal**:

```
publish: Internal
```

By setting this option, you create an internal Ingress Controller and a private load balancer.

3. Make any other modifications to the **install-config.yaml** file that you require. For more information about the parameters, see "Installation configuration parameters".
4. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

Additional resources

- [Installation configuration parameters for GCP](#)

9.12.6.3. Enabling Shielded VMs

You can use Shielded VMs when installing your cluster. Shielded VMs have extra security features including secure boot, firmware and integrity monitoring, and rootkit detection. For more information, see Google's documentation on [Shielded VMs](#).



NOTE

Shielded VMs are currently not supported on clusters with 64-bit ARM infrastructures.

Prerequisites

- You have created an **install-config.yaml** file.

Procedure

- Use a text editor to edit the **install-config.yaml** file prior to deploying your cluster and add one of the following stanzas:
 - a. To use shielded VMs for only control plane machines:

```
controlPlane:
  platform:
    gcp:
      secureBoot: Enabled
```

- b. To use shielded VMs for only compute machines:


```
compute:
  - platform:
      gcp:
        secureBoot: Enabled
```

- c. To use shielded VMs for all machines:

```
platform:
  gcp:
    defaultMachinePlatform:
      secureBoot: Enabled
```

9.12.6.4. Enabling Confidential VMs

You can use Confidential VMs when installing your cluster. Confidential VMs encrypt data while it is being processed. For more information, see Google's documentation on [Confidential Computing](#). You can enable Confidential VMs and Shielded VMs at the same time, although they are not dependent on each other.



NOTE

Confidential VMs are currently not supported on 64-bit ARM architectures.

Prerequisites

- You have created an **install-config.yaml** file.

Procedure

- Use a text editor to edit the **install-config.yaml** file prior to deploying your cluster and add one of the following stanzas:
 - a. To use confidential VMs for only control plane machines:

```
controlPlane:
  platform:
    gcp:
      confidentialCompute: Enabled 1
      type: n2d-standard-8 2
      onHostMaintenance: Terminate 3
```

- 1 Enable confidential VMs.
- 2 Specify a machine type that supports Confidential VMs. Confidential VMs require the N2D or C2D series of machine types. For more information on supported machine types, see [Supported operating systems and machine types](#).
- 3 Specify the behavior of the VM during a host maintenance event, such as a hardware or software update. For a machine that uses Confidential VM, this value must be set to **Terminate**, which stops the VM. Confidential VMs do not support live VM migration.

b. To use confidential VMs for only compute machines:

```
compute:
- platform:
  gcp:
    confidentialCompute: Enabled
    type: n2d-standard-8
    onHostMaintenance: Terminate
```

c. To use confidential VMs for all machines:


```
platform:
  gcp:
    defaultMachinePlatform:
      confidentialCompute: Enabled
      type: n2d-standard-8
      onHostMaintenance: Terminate
```

9.12.6.5. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```

apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5


```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.



NOTE

The installation program does not support the proxy **readinessEndpoints** field.



NOTE

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.



NOTE

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

9.12.6.6. Creating the Kubernetes manifest and Ignition config files

Because you must modify some cluster definition files and manually start the cluster machines, you must generate the Kubernetes manifest and Ignition config files that the cluster needs to configure the machines.

The installation configuration file transforms into the Kubernetes manifests. The manifests wrap into the Ignition configuration files, which are later used to configure the cluster machines.



IMPORTANT

- The Ignition config files that the OpenShift Container Platform installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

Prerequisites

- You obtained the OpenShift Container Platform installation program. For a restricted network installation, these files are on your mirror host.
- You created the **install-config.yaml** installation configuration file.

Procedure

1. Change to the directory that contains the OpenShift Container Platform installation program and generate the Kubernetes manifests for the cluster:

```
$ ./openshift-install create manifests --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the installation directory that contains the **install-config.yaml** file you created.

2. Remove the Kubernetes manifest files that define the control plane machines:

```
$ rm -f <installation_directory>/openshift/99_openshift-cluster-api_master-machines-*.yaml
```


By removing these files, you prevent the cluster from automatically generating control plane machines.

- Remove the Kubernetes manifest files that define the control plane machine set:

```
$ rm -f <installation_directory>/openshift/99_openshift-machine-api_master-control-plane-machine-set.yaml
```

- Optional: If you do not want the cluster to provision compute machines, remove the Kubernetes manifest files that define the worker machines:

```
$ rm -f <installation_directory>/openshift/99_openshift-cluster-api_worker-machineset-*.yaml
```



IMPORTANT

If you disabled the **MachineAPI** capability when installing a cluster on user-provisioned infrastructure, you must remove the Kubernetes manifest files that define the worker machines. Otherwise, your cluster fails to install.

Because you create and manage the worker machines yourself, you do not need to initialize these machines.

- Check that the **mastersSchedulable** parameter in the `<installation_directory>/manifests/cluster-scheduler-02-config.yml` Kubernetes manifest file is set to **false**. This setting prevents pods from being scheduled on the control plane machines:
 - Open the `<installation_directory>/manifests/cluster-scheduler-02-config.yml` file.
 - Locate the **mastersSchedulable** parameter and ensure that it is set to **false**.
 - Save and exit the file.
- Optional: If you do not want [the Ingress Operator](#) to create DNS records on your behalf, remove the **privateZone** and **publicZone** sections from the `<installation_directory>/manifests/cluster-dns-02-config.yml` DNS configuration file:

```
apiVersion: config.openshift.io/v1
kind: DNS
metadata:
  creationTimestamp: null
  name: cluster
spec:
  baseDomain: example.openshift.com
  privateZone: 1
    id: mycluster-100419-private-zone
  publicZone: 2
    id: example.openshift.com
status: {}
```

- Remove this section completely.


If you do so, you must add ingress DNS records manually in a later step.

7. To create the Ignition configuration files, run the following command from the directory that contains the installation program:

```
$ ./openshift-install create ignition-configs --dir <installation_directory> 1
```

- 1 For **<installation_directory>**, specify the same installation directory.

Ignition config files are created for the bootstrap, control plane, and compute nodes in the installation directory. The **kubeadmin-password** and **kubeconfig** files are created in the **./<installation_directory>/auth** directory:



Additional resources

- [Optional: Adding the ingress DNS records](#)

9.12.7. Exporting common variables

9.12.7.1. Extracting the infrastructure name

The Ignition config files contain a unique cluster identifier that you can use to uniquely identify your cluster in Google Cloud Platform (GCP). The infrastructure name is also used to locate the appropriate GCP resources during an OpenShift Container Platform installation. The provided Deployment Manager templates contain references to this infrastructure name, so you must extract it.

Prerequisites

- You obtained the OpenShift Container Platform installation program and the pull secret for your cluster.
- You generated the Ignition config files for your cluster.
- You installed the **jq** package.

Procedure

- To extract and view the infrastructure name from the Ignition config file metadata, run the following command:

```
$ jq -r .infraID <installation_directory>/metadata.json 1
```

- 1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.


Example output

```
openshift-vw9j6 1
```

- 1** The output of this command is your cluster name and a random string.

9.12.7.2. Exporting common variables for Deployment Manager templates

You must export a common set of variables that are used with the provided Deployment Manager templates used to assist in completing a user-provided infrastructure install on Google Cloud Platform (GCP).



NOTE

Specific Deployment Manager templates can also require additional exported variables, which are detailed in their related procedures.

Prerequisites

- Obtain the OpenShift Container Platform installation program and the pull secret for your cluster.
- Generate the Ignition config files for your cluster.
- Install the **jq** package.

Procedure

1. Export the following common variables to be used by the provided Deployment Manager templates:


```
$ export BASE_DOMAIN='<base_domain>'
$ export BASE_DOMAIN_ZONE_NAME='<base_domain_zone_name>'
$ export NETWORK_CIDR='10.0.0.0/16'
$ export MASTER_SUBNET_CIDR='10.0.0.0/17'
$ export WORKER_SUBNET_CIDR='10.0.128.0/17'

$ export KUBECONFIG='<installation_directory>/auth/kubeconfig' 1
$ export CLUSTER_NAME=`jq -r .clusterName <installation_directory>/metadata.json`
$ export INFRA_ID=`jq -r .infraID <installation_directory>/metadata.json`
$ export PROJECT_NAME=`jq -r .gcp.projectID <installation_directory>/metadata.json`
$ export REGION=`jq -r .gcp.region <installation_directory>/metadata.json`
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

9.12.8. Creating a VPC in GCP

You must create a VPC in Google Cloud Platform (GCP) for your OpenShift Container Platform cluster to use. You can customize the VPC to meet your requirements. One way to create the VPC is to modify the provided Deployment Manager template.



NOTE

If you do not use the provided Deployment Manager template to create your GCP infrastructure, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- Configure a GCP account.
- Generate the Ignition config files for your cluster.

Procedure

1. Copy the template from the **Deployment Manager template for the VPC** section of this topic and save it as **01_vpc.py** on your computer. This template describes the VPC that your cluster requires.
2. Create a **01_vpc.yaml** resource definition file:

```
$ cat <<EOF >01_vpc.yaml
imports:
- path: 01_vpc.py

resources:
- name: cluster-vpc
  type: 01_vpc.py
  properties:
    infra_id: '${INFRA_ID}' 1
    region: '${REGION}' 2
    master_subnet_cidr: '${MASTER_SUBNET_CIDR}' 3
    worker_subnet_cidr: '${WORKER_SUBNET_CIDR}' 4
EOF
```

- 1** **infra_id** is the **INFRA_ID** infrastructure name from the extraction step.
- 2** **region** is the region to deploy the cluster into, for example **us-central1**.
- 3** **master_subnet_cidr** is the CIDR for the master subnet, for example **10.0.0.0/17**.
- 4** **worker_subnet_cidr** is the CIDR for the worker subnet, for example **10.0.128.0/17**.

3. Create the deployment by using the **gcloud** CLI:

```
$ gcloud deployment-manager deployments create ${INFRA_ID}-vpc --config 01_vpc.yaml
```

9.12.8.1. Deployment Manager template for the VPC

You can use the following Deployment Manager template to deploy the VPC that you need for your OpenShift Container Platform cluster:

Example 9.117. 01_vpc.py Deployment Manager template

–

```

def GenerateConfig(context):

    resources = [{
        'name': context.properties['infra_id'] + '-network',
        'type': 'compute.v1.network',
        'properties': {
            'region': context.properties['region'],
            'autoCreateSubnetworks': False
        }
    }, {
        'name': context.properties['infra_id'] + '-master-subnet',
        'type': 'compute.v1.subnetwork',
        'properties': {
            'region': context.properties['region'],
            'network': '$(ref.' + context.properties['infra_id'] + '-network.selfLink)',
            'ipCidrRange': context.properties['master_subnet_cidr']
        }
    }, {
        'name': context.properties['infra_id'] + '-worker-subnet',
        'type': 'compute.v1.subnetwork',
        'properties': {
            'region': context.properties['region'],
            'network': '$(ref.' + context.properties['infra_id'] + '-network.selfLink)',
            'ipCidrRange': context.properties['worker_subnet_cidr']
        }
    }, {
        'name': context.properties['infra_id'] + '-router',
        'type': 'compute.v1.router',
        'properties': {
            'region': context.properties['region'],
            'network': '$(ref.' + context.properties['infra_id'] + '-network.selfLink)',
            'nats': [{
                'name': context.properties['infra_id'] + '-nat-master',
                'natIpAllocateOption': 'AUTO_ONLY',
                'minPortsPerVm': 7168,
                'sourceSubnetworkIpRangesToNat': 'LIST_OF_SUBNETWORKS',
                'subnetworks': [{
                    'name': '$(ref.' + context.properties['infra_id'] + '-master-subnet.selfLink)',
                    'sourceIpRangesToNat': ['ALL_IP_RANGES']
                }
            ]
        }
    }, {
        'name': context.properties['infra_id'] + '-nat-worker',
        'natIpAllocateOption': 'AUTO_ONLY',
        'minPortsPerVm': 512,
        'sourceSubnetworkIpRangesToNat': 'LIST_OF_SUBNETWORKS',
        'subnetworks': [{
            'name': '$(ref.' + context.properties['infra_id'] + '-worker-subnet.selfLink)',
            'sourceIpRangesToNat': ['ALL_IP_RANGES']
        }
    ]
    }
    ]

    return {'resources': resources}

```

9.12.9. Networking requirements for user-provisioned infrastructure

All the Red Hat Enterprise Linux CoreOS (RHCOS) machines require networking to be configured in **initramfs** during boot to fetch their Ignition config files.

9.12.9.1. Setting the cluster node hostnames through DHCP

On Red Hat Enterprise Linux CoreOS (RHCOS) machines, the hostname is set through NetworkManager. By default, the machines obtain their hostname through DHCP. If the hostname is not provided by DHCP, set statically through kernel arguments, or another method, it is obtained through a reverse DNS lookup. Reverse DNS lookup occurs after the network has been initialized on a node and can take time to resolve. Other system services can start prior to this and detect the hostname as **localhost** or similar. You can avoid this by using DHCP to provide the hostname for each cluster node.

Additionally, setting the hostnames through DHCP can bypass any manual DNS record name configuration errors in environments that have a DNS split-horizon implementation.

9.12.9.2. Network connectivity requirements

You must configure the network connectivity between machines to allow OpenShift Container Platform cluster components to communicate. Each machine must be able to resolve the hostnames of all other machines in the cluster.

This section provides details about the ports that are required.

Table 9.45. Ports used for all-machine to all-machine communications

Protocol	Port	Description
ICMP	N/A	Network reachability tests
TCP	1936	Metrics
	9000-9999	Host level services, including the node exporter on ports 9100-9101 and the Cluster Version Operator on port 9099 .
	10250-10259	The default ports that Kubernetes reserves
UDP	4789	VXLAN
	6081	Geneve
	9000-9999	Host level services, including the node exporter on ports 9100-9101 .
	500	IPsec IKE packets
	4500	IPsec NAT-T packets

Protocol	Port	Description
	123	Network Time Protocol (NTP) on UDP port 123 If an external NTP time server is configured, you must open UDP port 123 .
TCP/UDP	30000-32767	Kubernetes node port
ESP	N/A	IPsec Encapsulating Security Payload (ESP)

Table 9.46. Ports used for all-machine to control plane communications

Protocol	Port	Description
TCP	6443	Kubernetes API

Table 9.47. Ports used for control plane machine to control plane machine communications

Protocol	Port	Description
TCP	2379-2380	etcd server and peer ports

9.12.10. Creating load balancers in GCP

You must configure load balancers in Google Cloud Platform (GCP) for your OpenShift Container Platform cluster to use. One way to create these components is to modify the provided Deployment Manager template.



NOTE

If you do not use the provided Deployment Manager template to create your GCP infrastructure, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- Configure a GCP account.
- Generate the Ignition config files for your cluster.
- Create and configure a VPC and associated subnets in GCP.

Procedure

1. Copy the template from the **Deployment Manager template for the internal load balancer** section of this topic and save it as **02_lb_int.py** on your computer. This template describes the internal load balancing objects that your cluster requires.

2. For an external cluster, also copy the template from the **Deployment Manager template for the external load balancer** section of this topic and save it as **02_lb_ext.py** on your computer. This template describes the external load balancing objects that your cluster requires.
3. Export the variables that the deployment template uses:

- a. Export the cluster network location:

```
$ export CLUSTER_NETWORK=(`gcloud compute networks describe ${INFRA_ID}-network --format json | jq -r .selfLink`)
```

- b. Export the control plane subnet location:

```
$ export CONTROL_SUBNET=(`gcloud compute networks subnets describe ${INFRA_ID}-master-subnet --region=${REGION} --format json | jq -r .selfLink`)
```

- c. Export the three zones that the cluster uses:

```
$ export ZONE_0=(`gcloud compute regions describe ${REGION} --format=json | jq -r .zones[0] | cut -d "/" -f9`)
```

```
$ export ZONE_1=(`gcloud compute regions describe ${REGION} --format=json | jq -r .zones[1] | cut -d "/" -f9`)
```

```
$ export ZONE_2=(`gcloud compute regions describe ${REGION} --format=json | jq -r .zones[2] | cut -d "/" -f9`)
```

4. Create a **02_infra.yaml** resource definition file:

```
$ cat <<EOF >02_infra.yaml
imports:
- path: 02_lb_ext.py
- path: 02_lb_int.py 1
resources:
- name: cluster-lb-ext 2
  type: 02_lb_ext.py
  properties:
    infra_id: '${INFRA_ID}' 3
    region: '${REGION}' 4
- name: cluster-lb-int
  type: 02_lb_int.py
  properties:
    cluster_network: '${CLUSTER_NETWORK}'
    control_subnet: '${CONTROL_SUBNET}' 5
    infra_id: '${INFRA_ID}'
    region: '${REGION}'
    zones: 6
    - '${ZONE_0}'
    - '${ZONE_1}'
    - '${ZONE_2}'
EOF
```

1 2 Required only when deploying an external cluster.

- 3 **infra_id** is the **INFRA_ID** infrastructure name from the extraction step.
- 4 **region** is the region to deploy the cluster into, for example **us-central1**.
- 5 **control_subnet** is the URI to the control subnet.
- 6 **zones** are the zones to deploy the control plane instances into, like **us-east1-b**, **us-east1-c**, and **us-east1-d**.

5. Create the deployment by using the **gcloud** CLI:

```
$ gcloud deployment-manager deployments create ${INFRA_ID}-infra --config 02_infra.yaml
```

6. Export the cluster IP address:

```
$ export CLUSTER_IP=(`gcloud compute addresses describe ${INFRA_ID}-cluster-ip --region=${REGION} --format json | jq -r .address`)
```

7. For an external cluster, also export the cluster public IP address:

```
$ export CLUSTER_PUBLIC_IP=(`gcloud compute addresses describe ${INFRA_ID}-cluster-public-ip --region=${REGION} --format json | jq -r .address`)
```

9.12.10.1. Deployment Manager template for the external load balancer

You can use the following Deployment Manager template to deploy the external load balancer that you need for your OpenShift Container Platform cluster:

Example 9.118. 02_lb_ext.py Deployment Manager template

```
def GenerateConfig(context):

    resources = [{
        'name': context.properties['infra_id'] + '-cluster-public-ip',
        'type': 'compute.v1.address',
        'properties': {
            'region': context.properties['region']
        }
    }, {
        # Refer to docs/dev/kube-apiserver-health-check.md on how to correctly setup health check
        # probe for kube-apiserver
        'name': context.properties['infra_id'] + '-api-http-health-check',
        'type': 'compute.v1.httpHealthCheck',
        'properties': {
            'port': 6080,
            'requestPath': '/readyz'
        }
    }, {
        'name': context.properties['infra_id'] + '-api-target-pool',
        'type': 'compute.v1.targetPool',
        'properties': {
            'region': context.properties['region'],
            'healthChecks': ['$ (ref.' + context.properties['infra_id'] + '-api-http-health-check.selfLink)'],
```

```

        'instances': []
    }
}, {
    'name': context.properties['infra_id'] + '-api-forwarding-rule',
    'type': 'compute.v1.forwardingRule',
    'properties': {
        'region': context.properties['region'],
        'IPAddress': '${ref.' + context.properties['infra_id'] + '-cluster-public-ip.selfLink}',
        'target': '${ref.' + context.properties['infra_id'] + '-api-target-pool.selfLink}',
        'portRange': '6443'
    }
}
]]

return {'resources': resources}

```

9.12.10.2. Deployment Manager template for the internal load balancer

You can use the following Deployment Manager template to deploy the internal load balancer that you need for your OpenShift Container Platform cluster:

Example 9.119. 02_lb_int.py Deployment Manager template

```

def GenerateConfig(context):

    backends = []
    for zone in context.properties['zones']:
        backends.append({
            'group': '${ref.' + context.properties['infra_id'] + '-master-' + zone + '-ig' + '.selfLink}'
        })

    resources = [{
        'name': context.properties['infra_id'] + '-cluster-ip',
        'type': 'compute.v1.address',
        'properties': {
            'addressType': 'INTERNAL',
            'region': context.properties['region'],
            'subnetwork': context.properties['control_subnet']
        }
    }, {
        # Refer to docs/dev/kube-apiserver-health-check.md on how to correctly setup health check
        # probe for kube-apiserver
        'name': context.properties['infra_id'] + '-api-internal-health-check',
        'type': 'compute.v1.healthCheck',
        'properties': {
            'httpsHealthCheck': {
                'port': 6443,
                'requestPath': '/readyz'
            },
            'type': "HTTPS"
        }
    }, {
        'name': context.properties['infra_id'] + '-api-internal-backend-service',
        'type': 'compute.v1.regionBackendService',
        'properties': {

```



```

        'backends': backends,
        'healthChecks': ['$$(ref.' + context.properties['infra_id'] + '-api-internal-health-
check.selfLink)'],
        'loadBalancingScheme': 'INTERNAL',
        'region': context.properties['region'],
        'protocol': 'TCP',
        'timeoutSec': 120
    }
}, {
    'name': context.properties['infra_id'] + '-api-internal-forwarding-rule',
    'type': 'compute.v1.forwardingRule',
    'properties': {
        'backendService': '$$(ref.' + context.properties['infra_id'] + '-api-internal-backend-
service.selfLink)',
        'IPAddress': '$$(ref.' + context.properties['infra_id'] + '-cluster-ip.selfLink)',
        'loadBalancingScheme': 'INTERNAL',
        'ports': ['6443','22623'],
        'region': context.properties['region'],
        'subnetwork': context.properties['control_subnet']
    }
}
]]

for zone in context.properties['zones']:
    resources.append({
        'name': context.properties['infra_id'] + '-master-' + zone + '-ig',
        'type': 'compute.v1.instanceGroup',
        'properties': {
            'namedPorts': [
                {
                    'name': 'ignition',
                    'port': 22623
                }, {
                    'name': 'https',
                    'port': 6443
                }
            ],
            'network': context.properties['cluster_network'],
            'zone': zone
        }
    })

return {'resources': resources}

```

You will need this template in addition to the **02_lb_ext.py** template when you create an external cluster.

9.12.11. Creating a private DNS zone in GCP

You must configure a private DNS zone in Google Cloud Platform (GCP) for your OpenShift Container Platform cluster to use. One way to create this component is to modify the provided Deployment Manager template.



NOTE

If you do not use the provided Deployment Manager template to create your GCP infrastructure, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- Configure a GCP account.
- Generate the Ignition config files for your cluster.
- Create and configure a VPC and associated subnets in GCP.

Procedure

1. Copy the template from the **Deployment Manager template for the private DNS** section of this topic and save it as **02_dns.py** on your computer. This template describes the private DNS objects that your cluster requires.
2. Create a **02_dns.yaml** resource definition file:

```
$ cat <<EOF >02_dns.yaml
imports:
- path: 02_dns.py

resources:
- name: cluster-dns
  type: 02_dns.py
  properties:
    infra_id: '${INFRA_ID}' 1
    cluster_domain: '${CLUSTER_NAME}.${BASE_DOMAIN}' 2
    cluster_network: '${CLUSTER_NETWORK}' 3
EOF
```

- 1** **infra_id** is the **INFRA_ID** infrastructure name from the extraction step.
- 2** **cluster_domain** is the domain for the cluster, for example **openshift.example.com**.
- 3** **cluster_network** is the **selfLink** URL to the cluster network.

3. Create the deployment by using the **gcloud** CLI:

```
$ gcloud deployment-manager deployments create ${INFRA_ID}-dns --config 02_dns.yaml
```

4. The templates do not create DNS entries due to limitations of Deployment Manager, so you must create them manually:
 - a. Add the internal DNS entries:

```
$ if [ -f transaction.yaml ]; then rm transaction.yaml; fi
$ gcloud dns record-sets transaction start --zone ${INFRA_ID}-private-zone
$ gcloud dns record-sets transaction add ${CLUSTER_IP} --name
```

```

api.${CLUSTER_NAME}.${BASE_DOMAIN}. --ttl 60 --type A --zone ${INFRA_ID}-
private-zone
$ gcloud dns record-sets transaction add ${CLUSTER_IP} --name api-
int.${CLUSTER_NAME}.${BASE_DOMAIN}. --ttl 60 --type A --zone ${INFRA_ID}-
private-zone
$ gcloud dns record-sets transaction execute --zone ${INFRA_ID}-private-zone

```

- b. For an external cluster, also add the external DNS entries:

```

$ if [ -f transaction.yaml ]; then rm transaction.yaml; fi
$ gcloud dns record-sets transaction start --zone ${BASE_DOMAIN_ZONE_NAME}
$ gcloud dns record-sets transaction add ${CLUSTER_PUBLIC_IP} --name
api.${CLUSTER_NAME}.${BASE_DOMAIN}. --ttl 60 --type A --zone
${BASE_DOMAIN_ZONE_NAME}
$ gcloud dns record-sets transaction execute --zone ${BASE_DOMAIN_ZONE_NAME}

```

9.12.11.1. Deployment Manager template for the private DNS

You can use the following Deployment Manager template to deploy the private DNS that you need for your OpenShift Container Platform cluster:

Example 9.120. `02_dns.py` Deployment Manager template

```

def GenerateConfig(context):


    resources = [{
        'name': context.properties['infra_id'] + '-private-zone',
        'type': 'dns.v1.managedZone',
        'properties': {
            'description': "",
            'dnsName': context.properties['cluster_domain'] + '.',
            'visibility': 'private',
            'privateVisibilityConfig': {
                'networks': [{
                    'networkUrl': context.properties['cluster_network']
                }]
            }
        }
    }]

    return {'resources': resources}

```

9.12.12. Creating firewall rules in GCP

You must create firewall rules in Google Cloud Platform (GCP) for your OpenShift Container Platform cluster to use. One way to create these components is to modify the provided Deployment Manager template.



NOTE

If you do not use the provided Deployment Manager template to create your GCP infrastructure, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- Configure a GCP account.
- Generate the Ignition config files for your cluster.
- Create and configure a VPC and associated subnets in GCP.

Procedure

1. Copy the template from the **Deployment Manager template for firewall rules** section of this topic and save it as **03_firewall.py** on your computer. This template describes the security groups that your cluster requires.
2. Create a **03_firewall.yaml** resource definition file:

```
$ cat <<EOF >03_firewall.yaml
imports:
- path: 03_firewall.py

resources:
- name: cluster-firewall
  type: 03_firewall.py
  properties:
    allowed_external_cidr: '0.0.0.0/0' 1
    infra_id: '${INFRA_ID}' 2
    cluster_network: '${CLUSTER_NETWORK}' 3
    network_cidr: '${NETWORK_CIDR}' 4
EOF
```

- 1 **allowed_external_cidr** is the CIDR range that can access the cluster API and SSH to the bootstrap host. For an internal cluster, set this value to **\${NETWORK_CIDR}**.
- 2 **infra_id** is the **INFRA_ID** infrastructure name from the extraction step.
- 3 **cluster_network** is the **selfLink** URL to the cluster network.
- 4 **network_cidr** is the CIDR of the VPC network, for example **10.0.0.0/16**.

3. Create the deployment by using the **gcloud** CLI:

```
$ gcloud deployment-manager deployments create ${INFRA_ID}-firewall --config
03_firewall.yaml
```

9.12.12.1. Deployment Manager template for firewall rules

You can use the following Deployment Manager template to deploy the firewall rules that you need for your OpenShift Container Platform cluster:

Example 9.121. 03_firewall.py Deployment Manager template

```
def GenerateConfig(context):

    resources = [{
        'name': context.properties['infra_id'] + '-bootstrap-in-ssh',
        'type': 'compute.v1.firewall',
        'properties': {
            'network': context.properties['cluster_network'],
            'allowed': [{
                'IPProtocol': 'tcp',
                'ports': ['22']
            }],
            'sourceRanges': [context.properties['allowed_external_cidr']],
            'targetTags': [context.properties['infra_id'] + '-bootstrap']
        }
    ], {
        'name': context.properties['infra_id'] + '-api',
        'type': 'compute.v1.firewall',
        'properties': {
            'network': context.properties['cluster_network'],
            'allowed': [{
                'IPProtocol': 'tcp',
                'ports': ['6443']
            }],
            'sourceRanges': [context.properties['allowed_external_cidr']],
            'targetTags': [context.properties['infra_id'] + '-master']
        }
    ], {
        'name': context.properties['infra_id'] + '-health-checks',
        'type': 'compute.v1.firewall',
        'properties': {
            'network': context.properties['cluster_network'],
            'allowed': [{
                'IPProtocol': 'tcp',
                'ports': ['6080', '6443', '22624']
            }],
            'sourceRanges': ['35.191.0.0/16', '130.211.0.0/22', '209.85.152.0/22', '209.85.204.0/22'],
            'targetTags': [context.properties['infra_id'] + '-master']
        }
    ], {
        'name': context.properties['infra_id'] + '-etcd',
        'type': 'compute.v1.firewall',
        'properties': {
            'network': context.properties['cluster_network'],
            'allowed': [{
                'IPProtocol': 'tcp',
                'ports': ['2379-2380']
            }],
            'sourceTags': [context.properties['infra_id'] + '-master'],
            'targetTags': [context.properties['infra_id'] + '-master']
        }
    ], {
```

```
'name': context.properties['infra_id'] + '-control-plane',
'type': 'compute.v1.firewall',
'properties': {
  'network': context.properties['cluster_network'],
  'allowed': [{
    'IPProtocol': 'tcp',
    'ports': ['10257']
  },{
    'IPProtocol': 'tcp',
    'ports': ['10259']
  },{
    'IPProtocol': 'tcp',
    'ports': ['22623']
  }],
  'sourceTags': [
    context.properties['infra_id'] + '-master',
    context.properties['infra_id'] + '-worker'
  ],
  'targetTags': [context.properties['infra_id'] + '-master']
}
}, {
'name': context.properties['infra_id'] + '-internal-network',
'type': 'compute.v1.firewall',
'properties': {
  'network': context.properties['cluster_network'],
  'allowed': [{
    'IPProtocol': 'icmp'
  },{
    'IPProtocol': 'tcp',
    'ports': ['22']
  }],
  'sourceRanges': [context.properties['network_cidr']],
  'targetTags': [
    context.properties['infra_id'] + '-master',
    context.properties['infra_id'] + '-worker'
  ]
}
}, {
'name': context.properties['infra_id'] + '-internal-cluster',
'type': 'compute.v1.firewall',
'properties': {
  'network': context.properties['cluster_network'],
  'allowed': [{
    'IPProtocol': 'udp',
    'ports': ['4789', '6081']
  },{
    'IPProtocol': 'udp',
    'ports': ['500', '4500']
  },{
    'IPProtocol': 'esp',
  },{
    'IPProtocol': 'tcp',
    'ports': ['9000-9999']
  },{
    'IPProtocol': 'udp',
    'ports': ['9000-9999']
  }
}
```

```

    },{
      'IPProtocol': 'tcp',
      'ports': ['10250']
    },{
      'IPProtocol': 'tcp',
      'ports': ['30000-32767']
    },{
      'IPProtocol': 'udp',
      'ports': ['30000-32767']
    }],
    'sourceTags': [
      context.properties['infra_id'] + '-master',
      context.properties['infra_id'] + '-worker'
    ],
    'targetTags': [
      context.properties['infra_id'] + '-master',
      context.properties['infra_id'] + '-worker'
    ]
  }
}
}

return {'resources': resources}

```

9.12.13. Creating IAM roles in GCP

You must create IAM roles in Google Cloud Platform (GCP) for your OpenShift Container Platform cluster to use. One way to create these components is to modify the provided Deployment Manager template.



NOTE

If you do not use the provided Deployment Manager template to create your GCP infrastructure, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- Configure a GCP account.
- Generate the Ignition config files for your cluster.
- Create and configure a VPC and associated subnets in GCP.

Procedure

1. Copy the template from the **Deployment Manager template for IAM roles** section of this topic and save it as **03_iam.py** on your computer. This template describes the IAM roles that your cluster requires.
2. Create a **03_iam.yaml** resource definition file:

```
$ cat <<EOF >03_iam.yaml
```

```

imports:
- path: 03_iam.py
resources:
- name: cluster-iam
  type: 03_iam.py
properties:
  infra_id: '${INFRA_ID}' 1
EOF

```

- 1** `infra_id` is the `INFRA_ID` infrastructure name from the extraction step.

3. Create the deployment by using the **gcloud** CLI:

```
$ gcloud deployment-manager deployments create ${INFRA_ID}-iam --config 03_iam.yaml
```

4. Export the variable for the master service account:

```
$ export MASTER_SERVICE_ACCOUNT=$(gcloud iam service-accounts list --filter
"email~^${INFRA_ID}-m@${PROJECT_NAME}." --format json | jq -r '[0].email')
```

5. Export the variable for the worker service account:

```
$ export WORKER_SERVICE_ACCOUNT=$(gcloud iam service-accounts list --filter
"email~^${INFRA_ID}-w@${PROJECT_NAME}." --format json | jq -r '[0].email')
```

6. Export the variable for the subnet that hosts the compute machines:

```
$ export COMPUTE_SUBNET=$(gcloud compute networks subnets describe ${INFRA_ID}-
worker-subnet --region=${REGION} --format json | jq -r .selfLink)
```

7. The templates do not create the policy bindings due to limitations of Deployment Manager, so you must create them manually:

```

$ gcloud projects add-iam-policy-binding ${PROJECT_NAME} --member
"serviceAccount:${MASTER_SERVICE_ACCOUNT}" --role "roles/compute.instanceAdmin"
$ gcloud projects add-iam-policy-binding ${PROJECT_NAME} --member
"serviceAccount:${MASTER_SERVICE_ACCOUNT}" --role "roles/compute.networkAdmin"
$ gcloud projects add-iam-policy-binding ${PROJECT_NAME} --member
"serviceAccount:${MASTER_SERVICE_ACCOUNT}" --role "roles/compute.securityAdmin"
$ gcloud projects add-iam-policy-binding ${PROJECT_NAME} --member
"serviceAccount:${MASTER_SERVICE_ACCOUNT}" --role "roles/iam.serviceAccountUser"
$ gcloud projects add-iam-policy-binding ${PROJECT_NAME} --member
"serviceAccount:${MASTER_SERVICE_ACCOUNT}" --role "roles/storage.admin"

$ gcloud projects add-iam-policy-binding ${PROJECT_NAME} --member
"serviceAccount:${WORKER_SERVICE_ACCOUNT}" --role "roles/compute.viewer"
$ gcloud projects add-iam-policy-binding ${PROJECT_NAME} --member
"serviceAccount:${WORKER_SERVICE_ACCOUNT}" --role "roles/storage.admin"

```

8. Create a service account key and store it locally for later use:

```
$ gcloud iam service-accounts keys create service-account-key.json --iam-
account=${MASTER_SERVICE_ACCOUNT}
```


9.12.13.1. Deployment Manager template for IAM roles

You can use the following Deployment Manager template to deploy the IAM roles that you need for your OpenShift Container Platform cluster:

Example 9.122. 03_iam.py Deployment Manager template

```
def GenerateConfig(context):

    resources = [{
        'name': context.properties['infra_id'] + '-master-node-sa',
        'type': 'iam.v1.serviceAccount',
        'properties': {
            'accountId': context.properties['infra_id'] + '-m',
            'displayName': context.properties['infra_id'] + '-master-node'
        }
    }, {
        'name': context.properties['infra_id'] + '-worker-node-sa',
        'type': 'iam.v1.serviceAccount',
        'properties': {
            'accountId': context.properties['infra_id'] + '-w',
            'displayName': context.properties['infra_id'] + '-worker-node'
        }
    }
    ]

    return {'resources': resources}
```

9.12.14. Creating the RHCOS cluster image for the GCP infrastructure

You must use a valid Red Hat Enterprise Linux CoreOS (RHCOS) image for Google Cloud Platform (GCP) for your OpenShift Container Platform nodes.

Procedure

1. Obtain the RHCOS image from the [RHCOS image mirror](#) page.



IMPORTANT

The RHCOS images might not change with every release of OpenShift Container Platform. You must download an image with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Use the image version that matches your OpenShift Container Platform version if it is available.

The file name contains the OpenShift Container Platform version number in the format **rhcos-
<version>-<arch>-gcp.<arch>.tar.gz**.

2. Create the Google storage bucket:

```
$ gsutil mb gs://<bucket_name>
```

3. Upload the RHCOS image to the Google storage bucket:

```
$ gsutil cp <downloaded_image_file_path>/rhcos-<version>-x86_64-gcp.x86_64.tar.gz
gs://<bucket_name>
```

4. Export the uploaded RHCOS image location as a variable:

```
$ export IMAGE_SOURCE=gs://<bucket_name>/rhcos-<version>-x86_64-gcp.x86_64.tar.gz
```

5. Create the cluster image:

```
$ gcloud compute images create "${INFRA_ID}-rhcos-image" \
--source-uri="${IMAGE_SOURCE}"
```

9.12.15. Creating the bootstrap machine in GCP

You must create the bootstrap machine in Google Cloud Platform (GCP) to use during OpenShift Container Platform cluster initialization. One way to create this machine is to modify the provided Deployment Manager template.



NOTE

If you do not use the provided Deployment Manager template to create your bootstrap machine, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- Configure a GCP account.
- Generate the Ignition config files for your cluster.
- Create and configure a VPC and associated subnets in GCP.
- Create and configure networking and load balancers in GCP.
- Create control plane and compute roles.
- Ensure pyOpenSSL is installed.

Procedure

1. Copy the template from the **Deployment Manager template for the bootstrap machine** section of this topic and save it as **04_bootstrap.py** on your computer. This template describes the bootstrap machine that your cluster requires.
2. Export the location of the Red Hat Enterprise Linux CoreOS (RHCOS) image that the installation program requires:

```
$ export CLUSTER_IMAGE=(`gcloud compute images describe ${INFRA_ID}-rhcos-image --
format json | jq -r .selfLink`)
```

3. Create a bucket and upload the **bootstrap.ign** file:

```
$ gsutil mb gs://${INFRA_ID}-bootstrap-ignition
```

```
$ gsutil cp <installation_directory>/bootstrap.ign gs://${INFRA_ID}-bootstrap-ignition/
```

4. Create a signed URL for the bootstrap instance to use to access the Ignition config. Export the URL from the output as a variable:

```
$ export BOOTSTRAP_IGN=`gsutil signurl -d 1h service-account-key.json gs://${INFRA_ID}-bootstrap-ignition/bootstrap.ign | grep "^gs:" | awk '{print $5}'`
```

5. Create a **04_bootstrap.yaml** resource definition file:

```
$ cat <<EOF >04_bootstrap.yaml
imports:
- path: 04_bootstrap.py

resources:
- name: cluster-bootstrap
  type: 04_bootstrap.py
  properties:
    infra_id: '${INFRA_ID}' 1
    region: '${REGION}' 2
    zone: '${ZONE_0}' 3

    cluster_network: '${CLUSTER_NETWORK}' 4
    control_subnet: '${CONTROL_SUBNET}' 5
    image: '${CLUSTER_IMAGE}' 6
    machine_type: 'n1-standard-4' 7
    root_volume_size: '128' 8

    bootstrap_ign: '${BOOTSTRAP_IGN}' 9
EOF
```

- 1 **infra_id** is the **INFRA_ID** infrastructure name from the extraction step.
- 2 **region** is the region to deploy the cluster into, for example **us-central1**.
- 3 **zone** is the zone to deploy the bootstrap instance into, for example **us-central1-b**.
- 4 **cluster_network** is the **selfLink** URL to the cluster network.
- 5 **control_subnet** is the **selfLink** URL to the control subnet.
- 6 **image** is the **selfLink** URL to the RHCOS image.
- 7 **machine_type** is the machine type of the instance, for example **n1-standard-4**.
- 8 **root_volume_size** is the boot disk size for the bootstrap machine.
- 9 **bootstrap_ign** is the URL output when creating a signed URL.

6. Create the deployment by using the **gcloud** CLI:

```
$ gcloud deployment-manager deployments create ${INFRA_ID}-bootstrap --config
04_bootstrap.yaml
```

7. The templates do not manage load balancer membership due to limitations of Deployment Manager, so you must add the bootstrap machine manually.

- a. Add the bootstrap instance to the internal load balancer instance group:

```
$ gcloud compute instance-groups unmanaged add-instances \
  ${INFRA_ID}-bootstrap-ig --zone=${ZONE_0} --instances=${INFRA_ID}-bootstrap
```

- b. Add the bootstrap instance group to the internal load balancer backend service:

```
$ gcloud compute backend-services add-backend \
  ${INFRA_ID}-api-internal-backend-service --region=${REGION} --instance-
group=${INFRA_ID}-bootstrap-ig --instance-group-zone=${ZONE_0}
```

9.12.15.1. Deployment Manager template for the bootstrap machine

You can use the following Deployment Manager template to deploy the bootstrap machine that you need for your OpenShift Container Platform cluster:

Example 9.123. 04_bootstrap.py Deployment Manager template

```
def GenerateConfig(context):

    resources = [{
        'name': context.properties['infra_id'] + '-bootstrap-public-ip',
        'type': 'compute.v1.address',
        'properties': {
            'region': context.properties['region']
        }
    }, {
        'name': context.properties['infra_id'] + '-bootstrap',
        'type': 'compute.v1.instance',
        'properties': {
            'disks': [{
                'autoDelete': True,
                'boot': True,
                'initializeParams': {
                    'diskSizeGb': context.properties['root_volume_size'],
                    'sourceImage': context.properties['image']
                }
            }],
            'machineType': 'zones/' + context.properties['zone'] + '/machineTypes/' +
context.properties['machine_type'],
            'metadata': {
                'items': [{
                    'key': 'user-data',
                    'value': '{"ignition":{"config":{"replace":{"source":"' + context.properties['bootstrap_ign']
+ ""},"version":"3.2.0"}}}'
                }],
            },
            'networkInterfaces': [{
```

```

    'subnetwork': context.properties['control_subnet'],
    'accessConfigs': [{
      'natIP': '${ref.' + context.properties['infra_id'] + '-bootstrap-public-ip.address}'
    }]
  },
  'tags': {
    'items': [
      context.properties['infra_id'] + '-master',
      context.properties['infra_id'] + '-bootstrap'
    ]
  },
  'zone': context.properties['zone']
}
}, {
  'name': context.properties['infra_id'] + '-bootstrap-ig',
  'type': 'compute.v1.instanceGroup',
  'properties': {
    'namedPorts': [
      {
        'name': 'ignition',
        'port': 22623
      }, {
        'name': 'https',
        'port': 6443
      }
    ],
    'network': context.properties['cluster_network'],
    'zone': context.properties['zone']
  }
}
]]

return {'resources': resources}

```

9.12.16. Creating the control plane machines in GCP

You must create the control plane machines in Google Cloud Platform (GCP) for your cluster to use. One way to create these machines is to modify the provided Deployment Manager template.



NOTE

If you do not use the provided Deployment Manager template to create your control plane machines, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- Configure a GCP account.
- Generate the Ignition config files for your cluster.
- Create and configure a VPC and associated subnets in GCP.
- Create and configure networking and load balancers in GCP.

- Create control plane and compute roles.
- Create the bootstrap machine.

Procedure

1. Copy the template from the **Deployment Manager template for control plane machines** section of this topic and save it as **05_control_plane.py** on your computer. This template describes the control plane machines that your cluster requires.
2. Export the following variable required by the resource definition:

```
$ export MASTER_IGNITION=`cat <installation_directory>/master.ign`
```

3. Create a **05_control_plane.yaml** resource definition file:

```
$ cat <<EOF >05_control_plane.yaml
imports:
- path: 05_control_plane.py

resources:
- name: cluster-control-plane
  type: 05_control_plane.py
  properties:
    infra_id: '${INFRA_ID}' 1
    zones: 2
    - '${ZONE_0}'
    - '${ZONE_1}'
    - '${ZONE_2}'

    control_subnet: '${CONTROL_SUBNET}' 3
    image: '${CLUSTER_IMAGE}' 4
    machine_type: 'n1-standard-4' 5
    root_volume_size: '128'
    service_account_email: '${MASTER_SERVICE_ACCOUNT}' 6

    ignition: '${MASTER_IGNITION}' 7
EOF
```

- 1 **infra_id** is the **INFRA_ID** infrastructure name from the extraction step.
- 2 **zones** are the zones to deploy the control plane instances into, for example **us-central1-a**, **us-central1-b**, and **us-central1-c**.
- 3 **control_subnet** is the **selfLink** URL to the control subnet.
- 4 **image** is the **selfLink** URL to the RHCOS image.
- 5 **machine_type** is the machine type of the instance, for example **n1-standard-4**.
- 6 **service_account_email** is the email address for the master service account that you created.
- 7 **ignition** is the contents of the **master.ign** file.

4. Create the deployment by using the **gcloud** CLI:

```
$ gcloud deployment-manager deployments create ${INFRA_ID}-control-plane --config
05_control_plane.yaml
```

5. The templates do not manage load balancer membership due to limitations of Deployment Manager, so you must add the control plane machines manually.

- Run the following commands to add the control plane machines to the appropriate instance groups:

```
$ gcloud compute instance-groups unmanaged add-instances ${INFRA_ID}-master-
${ZONE_0}-ig --zone=${ZONE_0} --instances=${INFRA_ID}-master-0
```

```
$ gcloud compute instance-groups unmanaged add-instances ${INFRA_ID}-master-
${ZONE_1}-ig --zone=${ZONE_1} --instances=${INFRA_ID}-master-1
```

```
$ gcloud compute instance-groups unmanaged add-instances ${INFRA_ID}-master-
${ZONE_2}-ig --zone=${ZONE_2} --instances=${INFRA_ID}-master-2
```

- For an external cluster, you must also run the following commands to add the control plane machines to the target pools:

```
$ gcloud compute target-pools add-instances ${INFRA_ID}-api-target-pool --instances-
zone="${ZONE_0}" --instances=${INFRA_ID}-master-0
```

```
$ gcloud compute target-pools add-instances ${INFRA_ID}-api-target-pool --instances-
zone="${ZONE_1}" --instances=${INFRA_ID}-master-1
```

```
$ gcloud compute target-pools add-instances ${INFRA_ID}-api-target-pool --instances-
zone="${ZONE_2}" --instances=${INFRA_ID}-master-2
```

9.12.16.1. Deployment Manager template for control plane machines

You can use the following Deployment Manager template to deploy the control plane machines that you need for your OpenShift Container Platform cluster:

Example 9.124. 05_control_plane.py Deployment Manager template

```
def GenerateConfig(context):

    resources = [{
        'name': context.properties['infra_id'] + '-master-0',
        'type': 'compute.v1.instance',
        'properties': {
            'disks': [{
                'autoDelete': True,
                'boot': True,
                'initializeParams': {
                    'diskSizeGb': context.properties['root_volume_size'],
                    'diskType': 'zones/' + context.properties['zones'][0] + '/diskTypes/pd-ssd',
                    'sourceImage': context.properties['image']
```

```

    }
  }],
  'machineType': 'zones/' + context.properties['zones'][0] + '/machineTypes/' +
context.properties['machine_type'],
  'metadata': {
    'items': [{
      'key': 'user-data',
      'value': context.properties['ignition']
    }]
  },
  'networkInterfaces': [{
    'subnetwork': context.properties['control_subnet']
  }],
  'serviceAccounts': [{
    'email': context.properties['service_account_email'],
    'scopes': ['https://www.googleapis.com/auth/cloud-platform']
  }],
  'tags': {
    'items': [
      context.properties['infra_id'] + '-master',
    ]
  },
  'zone': context.properties['zones'][0]
}
}, {
  'name': context.properties['infra_id'] + '-master-1',
  'type': 'compute.v1.instance',
  'properties': {
    'disks': [{
      'autoDelete': True,
      'boot': True,
      'initializeParams': {
        'diskSizeGb': context.properties['root_volume_size'],
        'diskType': 'zones/' + context.properties['zones'][1] + '/diskTypes/pd-ssd',
        'sourceImage': context.properties['image']
      }
    }]
  },
  'machineType': 'zones/' + context.properties['zones'][1] + '/machineTypes/' +
context.properties['machine_type'],
  'metadata': {
    'items': [{
      'key': 'user-data',
      'value': context.properties['ignition']
    }]
  },
  'networkInterfaces': [{
    'subnetwork': context.properties['control_subnet']
  }],
  'serviceAccounts': [{
    'email': context.properties['service_account_email'],
    'scopes': ['https://www.googleapis.com/auth/cloud-platform']
  }],
  'tags': {
    'items': [
      context.properties['infra_id'] + '-master',
    ]
  }
}

```



```

    },
    'zone': context.properties['zones'][1]
  }
}, {
  'name': context.properties['infra_id'] + '-master-2',
  'type': 'compute.v1.instance',
  'properties': {
    'disks': [{
      'autoDelete': True,
      'boot': True,
      'initializeParams': {
        'diskSizeGb': context.properties['root_volume_size'],
        'diskType': 'zones/' + context.properties['zones'][2] + '/diskTypes/pd-ssd',
        'sourceImage': context.properties['image']
      }
    }
  ],
  'machineType': 'zones/' + context.properties['zones'][2] + '/machineTypes/' +
context.properties['machine_type'],
  'metadata': {
    'items': [{
      'key': 'user-data',
      'value': context.properties['ignition']
    }
  ]
},
  'networkInterfaces': [{
    'subnetwork': context.properties['control_subnet']
  }],
  'serviceAccounts': [{
    'email': context.properties['service_account_email'],
    'scopes': ['https://www.googleapis.com/auth/cloud-platform']
  }],
  'tags': {
    'items': [
      context.properties['infra_id'] + '-master',
    ]
  },
  'zone': context.properties['zones'][2]
}
]]

return {'resources': resources}

```

9.12.17. Wait for bootstrap completion and remove bootstrap resources in GCP

After you create all of the required infrastructure in Google Cloud Platform (GCP), wait for the bootstrap process to complete on the machines that you provisioned by using the Ignition config files that you generated with the installation program.

Prerequisites

- Configure a GCP account.
- Generate the Ignition config files for your cluster.

- Create and configure a VPC and associated subnets in GCP.
- Create and configure networking and load balancers in GCP.
- Create control plane and compute roles.
- Create the bootstrap machine.
- Create the control plane machines.

Procedure

1. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install wait-for bootstrap-complete --dir <installation_directory> \ 1
--log-level info 2
```

- 1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

- 2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

If the command exits without a **FATAL** warning, your production control plane has initialized.

2. Delete the bootstrap resources:

```
$ gcloud compute backend-services remove-backend ${INFRA_ID}-api-internal-backend-
service --region=${REGION} --instance-group=${INFRA_ID}-bootstrap-ig --instance-group-
zone=${ZONE_0}
```

```
$ gsutil rm gs://${INFRA_ID}-bootstrap-ignition/bootstrap.ign
```

```
$ gsutil rb gs://${INFRA_ID}-bootstrap-ignition
```

```
$ gcloud deployment-manager deployments delete ${INFRA_ID}-bootstrap
```

9.12.18. Creating additional worker machines in GCP

You can create worker machines in Google Cloud Platform (GCP) for your cluster to use by launching individual instances discretely or by automated processes outside the cluster, such as auto scaling groups. You can also take advantage of the built-in cluster scaling mechanisms and the machine API in OpenShift Container Platform.

In this example, you manually launch one instance by using the Deployment Manager template. Additional instances can be launched by including additional resources of type **06_worker.py** in the file.



NOTE

If you do not use the provided Deployment Manager template to create your worker machines, you must review the provided information and manually create the infrastructure. If your cluster does not initialize correctly, you might have to contact Red Hat support with your installation logs.

Prerequisites

- Configure a GCP account.
- Generate the Ignition config files for your cluster.
- Create and configure a VPC and associated subnets in GCP.
- Create and configure networking and load balancers in GCP.
- Create control plane and compute roles.
- Create the bootstrap machine.
- Create the control plane machines.

Procedure

1. Copy the template from the **Deployment Manager template for worker machines** section of this topic and save it as **06_worker.py** on your computer. This template describes the worker machines that your cluster requires.
2. Export the variables that the resource definition uses.

- a. Export the subnet that hosts the compute machines:

```
$ export COMPUTE_SUBNET=$(gcloud compute networks subnets describe
${INFRA_ID}-worker-subnet --region=${REGION} --format json | jq -r '.selfLink')
```

- b. Export the email address for your service account:

```
$ export WORKER_SERVICE_ACCOUNT=$(gcloud iam service-accounts list --filter
"email~^${INFRA_ID}-w@${PROJECT_NAME}." --format json | jq -r '.[0].email')
```

- c. Export the location of the compute machine Ignition config file:

```
$ export WORKER_IGNITION=`cat <installation_directory>/worker.ign`
```

3. Create a **06_worker.yaml** resource definition file:

```
$ cat <<EOF >06_worker.yaml
imports:
- path: 06_worker.py

resources:
- name: 'worker-0' 1
  type: 06_worker.py
  properties:
    infra_id: '${INFRA_ID}' 2
    zone: '${ZONE_0}' 3
    compute_subnet: '${COMPUTE_SUBNET}' 4
    image: '${CLUSTER_IMAGE}' 5
    machine_type: 'n1-standard-4' 6
    root_volume_size: '128'
```

```

    service_account_email: '${WORKER_SERVICE_ACCOUNT}' 7
    ignition: '${WORKER_IGNITION}' 8
- name: 'worker-1'
  type: 06_worker.py
  properties:
    infra_id: '${INFRA_ID}' 9
    zone: '${ZONE_1}' 10
    compute_subnet: '${COMPUTE_SUBNET}' 11
    image: '${CLUSTER_IMAGE}' 12
    machine_type: 'n1-standard-4' 13
    root_volume_size: '128'
    service_account_email: '${WORKER_SERVICE_ACCOUNT}' 14
    ignition: '${WORKER_IGNITION}' 15
EOF

```

- 1 **name** is the name of the worker machine, for example **worker-0**.
- 2 9 **infra_id** is the **INFRA_ID** infrastructure name from the extraction step.
- 3 10 **zone** is the zone to deploy the worker machine into, for example **us-central1-a**.
- 4 11 **compute_subnet** is the **selfLink** URL to the compute subnet.
- 5 12 **image** is the **selfLink** URL to the RHCOS image. ¹
- 6 13 **machine_type** is the machine type of the instance, for example **n1-standard-4**.
- 7 14 **service_account_email** is the email address for the worker service account that you created.
- 8 15 **ignition** is the contents of the **worker.ign** file.

4. Optional: If you want to launch additional instances, include additional resources of type **06_worker.py** in your **06_worker.yaml** resource definition file.
5. Create the deployment by using the **gcloud** CLI:

```

$ gcloud deployment-manager deployments create ${INFRA_ID}-worker --config
06_worker.yaml

```

1. To use a GCP Marketplace image, specify the offer to use:
 - OpenShift Container Platform: <https://www.googleapis.com/compute/v1/projects/redhat-marketplace-public/global/images/redhat-coreos-ocp-413-x86-64-202305021736>
 - OpenShift Platform Plus: <https://www.googleapis.com/compute/v1/projects/redhat-marketplace-public/global/images/redhat-coreos-opp-413-x86-64-202305021736>
 - OpenShift Kubernetes Engine: <https://www.googleapis.com/compute/v1/projects/redhat-marketplace-public/global/images/redhat-coreos-oke-413-x86-64-202305021736>

9.12.18.1. Deployment Manager template for worker machines

You can use the following Deployment Manager template to deploy the worker machines that you need for your OpenShift Container Platform cluster:

Example 9.125. `06_worker.py` Deployment Manager template

```
def GenerateConfig(context):

    resources = [{
        'name': context.properties['infra_id'] + '-' + context.env['name'],
        'type': 'compute.v1.instance',
        'properties': {
            'disks': [{
                'autoDelete': True,
                'boot': True,
                'initializeParams': {
                    'diskSizeGb': context.properties['root_volume_size'],
                    'sourceImage': context.properties['image']
                }
            }],
            'machineType': 'zones/' + context.properties['zone'] + '/machineTypes/' +
context.properties['machine_type'],
            'metadata': {
                'items': [{
                    'key': 'user-data',
                    'value': context.properties['ignition']
                }]
            },
            'networkInterfaces': [{
                'subnetwork': context.properties['compute_subnet']
            }],
            'serviceAccounts': [{
                'email': context.properties['service_account_email'],
                'scopes': ['https://www.googleapis.com/auth/cloud-platform']
            }],
            'tags': {
                'items': [
                    context.properties['infra_id'] + '-worker',
                ]
            },
            'zone': context.properties['zone']
        }
    ]

    return {'resources': resources}
```

9.12.19. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

9.12.20. Disabling the default OperatorHub catalog sources

Operator catalogs that source content provided by Red Hat and community projects are configured for OperatorHub by default during an OpenShift Container Platform installation. In a restricted network environment, you must disable the default catalogs as a cluster administrator.

Procedure

- Disable the sources for the default catalogs by adding **disableAllDefaultSources: true** to the **OperatorHub** object:

```
$ oc patch OperatorHub cluster --type json \  
-p '[{"op": "add", "path": "/spec/disableAllDefaultSources", "value": true}]'
```

TIP

Alternatively, you can use the web console to manage catalog sources. From the **Administration** → **Cluster Settings** → **Configuration** → **OperatorHub** page, click the **Sources** tab, where you can create, update, delete, disable, and enable individual sources.

9.12.21. Approving the certificate signing requests for your machines

When you add machines to a cluster, two pending certificate signing requests (CSRs) are generated for each machine that you added. You must confirm that these CSRs are approved or, if necessary, approve them yourself. The client requests must be approved first, followed by the server requests.

Prerequisites

- You added machines to your cluster.

Procedure

Procedure

1. Confirm that the cluster recognizes the machines:

```
$ oc get nodes
```

Example output

```
NAME      STATUS    ROLES    AGE    VERSION
master-0  Ready     master   63m    v1.29.4
master-1  Ready     master   63m    v1.29.4
master-2  Ready     master   64m    v1.29.4
```

The output lists all of the machines that you created.



NOTE

The preceding output might not include the compute nodes, also known as worker nodes, until some CSRs are approved.

2. Review the pending CSRs and ensure that you see the client requests with the **Pending** or **Approved** status for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

```
NAME      AGE    REQUESTOR                                     CONDITION
csr-8b2br  15m    system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper  Pending
csr-8vnps  15m    system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper  Pending
...
```

In this example, two machines are joining the cluster. You might see more approved CSRs in the list.

3. If the CSRs were not approved, after all of the pending CSRs for the machines you added are in **Pending** status, approve the CSRs for your cluster machines:



NOTE

Because the CSRs rotate automatically, approve your CSRs within an hour of adding the machines to the cluster. If you do not approve them within an hour, the certificates will rotate, and more than two certificates will be present for each node. You must approve all of these certificates. After the client CSR is approved, the Kubelet creates a secondary CSR for the serving certificate, which requires manual approval. Then, subsequent serving certificate renewal requests are automatically approved by the **machine-approver** if the Kubelet requests a new certificate with identical parameters.

**NOTE**

For clusters running on platforms that are not machine API enabled, such as bare metal and other user-provisioned infrastructure, you must implement a method of automatically approving the kubelet serving certificate requests (CSRs). If a request is not approved, then the **oc exec**, **oc rsh**, and **oc logs** commands cannot succeed, because a serving certificate is required when the API server connects to the kubelet. Any operation that contacts the Kubelet endpoint requires this certificate approval to be in place. The method must watch for new CSRs, confirm that the CSR was submitted by the **node-bootstrap** service account in the **system:node** or **system:admin** groups, and confirm the identity of the node.

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{"\n"}{{end}}{{end}}' | xargs --no-run-if-empty oc adm certificate approve
```

**NOTE**

Some Operators might not become available until some CSRs are approved.

- Now that your client requests are approved, you must review the server requests for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

```
NAME      AGE   REQUESTOR                                     CONDITION
csr-bfd72 5m26s system:node:ip-10-0-50-126.us-east-2.compute.internal
Pending
csr-c57lv 5m26s system:node:ip-10-0-95-157.us-east-2.compute.internal
Pending
...
```

- If the remaining CSRs are not approved, and are in the **Pending** status, approve the CSRs for your cluster machines:

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:


```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{"\n"}
{{end}}{{end}} | xargs oc adm certificate approve
```

6. After all client and server CSRs have been approved, the machines have the **Ready** status. Verify this by running the following command:

```
$ oc get nodes
```

Example output

```
NAME      STATUS  ROLES  AGE  VERSION
master-0  Ready   master 73m  v1.29.4
master-1  Ready   master 73m  v1.29.4
master-2  Ready   master 74m  v1.29.4
worker-0  Ready   worker 11m  v1.29.4
worker-1  Ready   worker 11m  v1.29.4
```



NOTE

It can take a few minutes after approval of the server CSRs for the machines to transition to the **Ready** status.

Additional information

- For more information on CSRs, see [Certificate Signing Requests](#).

9.12.22. Optional: Adding the ingress DNS records

If you removed the DNS zone configuration when creating Kubernetes manifests and generating Ignition configs, you must manually create DNS records that point at the ingress load balancer. You can create either a wildcard ***.apps.{baseDomain}**, or specific records. You can use A, CNAME, and other records per your requirements.

Prerequisites

- Configure a GCP account.
- Remove the DNS Zone configuration when creating Kubernetes manifests and generating Ignition configs.
- Create and configure a VPC and associated subnets in GCP.
- Create and configure networking and load balancers in GCP.
- Create control plane and compute roles.
- Create the bootstrap machine.
- Create the control plane machines.
- Create the worker machines.

Procedure

1. Wait for the Ingress router to create a load balancer and populate the **EXTERNAL-IP** field:

```
$ oc -n openshift-ingress get service router-default
```

Example output

```
NAME           TYPE           CLUSTER-IP   EXTERNAL-IP   PORT(S)           AGE
router-default LoadBalancer   172.30.18.154 35.233.157.184 80:32288/TCP,443:31215/TCP 98
```

2. Add the A record to your zones:

- To use A records:

- i. Export the variable for the router IP address:

```
$ export ROUTER_IP=`oc -n openshift-ingress get service router-default --no-headers | awk '{print $4}'`
```

- ii. Add the A record to the private zones:

```
$ if [ -f transaction.yaml ]; then rm transaction.yaml; fi
$ gcloud dns record-sets transaction start --zone ${INFRA_ID}-private-zone
$ gcloud dns record-sets transaction add ${ROUTER_IP} --name
  \*.apps.${CLUSTER_NAME}.${BASE_DOMAIN}. --ttl 300 --type A --zone
  ${INFRA_ID}-private-zone
$ gcloud dns record-sets transaction execute --zone ${INFRA_ID}-private-zone
```

- iii. For an external cluster, also add the A record to the public zones:

```
$ if [ -f transaction.yaml ]; then rm transaction.yaml; fi
$ gcloud dns record-sets transaction start --zone ${BASE_DOMAIN_ZONE_NAME}
$ gcloud dns record-sets transaction add ${ROUTER_IP} --name
  \*.apps.${CLUSTER_NAME}.${BASE_DOMAIN}. --ttl 300 --type A --zone
  ${BASE_DOMAIN_ZONE_NAME}
$ gcloud dns record-sets transaction execute --zone
  ${BASE_DOMAIN_ZONE_NAME}
```

- To add explicit domains instead of using a wildcard, create entries for each of the cluster's current routes:

```
$ oc get --all-namespaces -o jsonpath='{range .items[*]}{range .status.ingress[*]}{.host} {"\n"}{end}{end}' routes
```

Example output

```
oauth-openshift.apps.your.cluster.domain.example.com
console-openshift-console.apps.your.cluster.domain.example.com
downloads-openshift-console.apps.your.cluster.domain.example.com
alertmanager-main-openshift-monitoring.apps.your.cluster.domain.example.com
prometheus-k8s-openshift-monitoring.apps.your.cluster.domain.example.com
```

9.12.23. Completing a GCP installation on user-provisioned infrastructure

After you start the OpenShift Container Platform installation on Google Cloud Platform (GCP) user-provisioned infrastructure, you can monitor the cluster events until the cluster is ready.

Prerequisites

- Deploy the bootstrap machine for an OpenShift Container Platform cluster on user-provisioned GCP infrastructure.
- Install the **oc** CLI and log in.

Procedure


1. Complete the cluster installation:

```
$ ./openshift-install --dir <installation_directory> wait-for install-complete 1
```

Example output

```
INFO Waiting up to 30m0s for the cluster to initialize...
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

2. Observe the running state of your cluster.
 - a. Run the following command to view the current cluster version and status:

```
$ oc get clusterversion
```

Example output

```
NAME      VERSION  AVAILABLE  PROGRESSING  SINCE   STATUS
version  False    True        24m         Working towards 4.5.4: 99% complete
```

- b. Run the following command to view the Operators managed on the control plane by the Cluster Version Operator (CVO):

```
$ oc get clusteroperators
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.5.4	True	False	False	7m56s
cloud-credential	4.5.4	True	False	False	31m
cluster-autoscaler	4.5.4	True	False	False	16m
console	4.5.4	True	False	False	10m
csi-snapshot-controller	4.5.4	True	False	False	16m
dns	4.5.4	True	False	False	22m
etcd	4.5.4	False	False	False	25s
image-registry	4.5.4	True	False	False	16m
ingress	4.5.4	True	False	False	16m
insights	4.5.4	True	False	False	17m
kube-apiserver	4.5.4	True	False	False	19m
kube-controller-manager	4.5.4	True	False	False	20m
kube-scheduler	4.5.4	True	False	False	20m
kube-storage-version-migrator	4.5.4	True	False	False	16m
machine-api	4.5.4	True	False	False	22m
machine-config	4.5.4	True	False	False	22m
marketplace	4.5.4	True	False	False	16m
monitoring	4.5.4	True	False	False	10m
network	4.5.4	True	False	False	23m
node-tuning	4.5.4	True	False	False	23m
openshift-apiserver	4.5.4	True	False	False	17m
openshift-controller-manager	4.5.4	True	False	False	15m
openshift-samples	4.5.4	True	False	False	16m
operator-lifecycle-manager	4.5.4	True	False	False	22m
operator-lifecycle-manager-catalog	4.5.4	True	False	False	22m
operator-lifecycle-manager-packageserver	4.5.4	True	False	False	18m
service-ca	4.5.4	True	False	False	23m
service-catalog-apiserver	4.5.4	True	False	False	23m
service-catalog-controller-manager	4.5.4	True	False	False	23m
storage	4.5.4	True	False	False	17m

- c. Run the following command to view your cluster pods:

```
$ oc get pods --all-namespaces
```

Example output

NAMESPACE	READY	STATUS	RESTARTS	NAME	AGE
kube-system				etcd-member-ip-10-0-3-111.us-east-	
2.compute.internal	1/1	Running	0		35m
kube-system				etcd-member-ip-10-0-3-239.us-east-	
2.compute.internal	1/1	Running	0		37m
kube-system				etcd-member-ip-10-0-3-24.us-east-	
2.compute.internal	1/1	Running	0		35m

```

openshift-apiserver-operator          openshift-apiserver-operator-6d6674f4f4-
h7t2t          1/1    Running    1    37m
openshift-apiserver                    apiserver-fm48r
1/1    Running    0    30m
openshift-apiserver                    apiserver-fxkvv
1/1    Running    0    29m
openshift-apiserver                    apiserver-q85nm
1/1    Running    0    29m
...
openshift-service-ca-operator          openshift-service-ca-operator-66ff6dc6cd-
9r257          1/1    Running    0    37m
openshift-service-ca                    apiservice-cabundle-injector-695b6bcbc-cl5hm
1/1    Running    0    35m
openshift-service-ca                    configmap-cabundle-injector-8498544d7-
25qn6          1/1    Running    0    35m
openshift-service-ca                    service-serving-cert-signer-6445fc9c6-wqdqn
1/1    Running    0    35m
openshift-service-catalog-apiserver-operator  openshift-service-catalog-apiserver-
operator-549f44668b-b5q2w  1/1    Running    0    32m
openshift-service-catalog-controller-manager-operator  openshift-service-catalog-
controller-manager-operator-b78cr2lnm  1/1    Running    0    31m

```

When the current cluster version is **AVAILABLE**, the installation is complete.

9.12.24. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service

9.12.25. Next steps

- [Customize your cluster](#).
- [Configure image streams](#) for the Cluster Samples Operator and the **must-gather** tool.
- Learn how to [use Operator Lifecycle Manager \(OLM\) on restricted networks](#).
- If the mirror registry that you used to install your cluster has a trusted CA, add it to the cluster by [configuring additional trust stores](#).
- If necessary, you can [opt out of remote health reporting](#).
- If necessary, see [Registering your disconnected cluster](#)

9.13. INSTALLING A THREE-NODE CLUSTER ON GCP

In OpenShift Container Platform version 4.16, you can install a three-node cluster on Google Cloud Platform (GCP). A three-node cluster consists of three control plane machines, which also act as compute machines. This type of cluster provides a smaller, more resource efficient cluster, for cluster administrators and developers to use for testing, development, and production.

You can install a three-node cluster using either installer-provisioned or user-provisioned infrastructure.

9.13.1. Configuring a three-node cluster

You configure a three-node cluster by setting the number of worker nodes to **0** in the **install-config.yaml** file before deploying the cluster. Setting the number of worker nodes to **0** ensures that the control plane machines are schedulable. This allows application workloads to be scheduled to run from the control plane nodes.



NOTE

Because application workloads run from control plane nodes, additional subscriptions are required, as the control plane nodes are considered to be compute nodes.

Prerequisites

- You have an existing **install-config.yaml** file.

Procedure

1. Set the number of compute replicas to **0** in your **install-config.yaml** file, as shown in the following **compute** stanza:

Example **install-config.yaml** file for a three-node cluster

```
apiVersion: v1
baseDomain: example.com
compute:
- name: worker
  platform: {}
  replicas: 0
# ...
```

2. If you are deploying a cluster with user-provisioned infrastructure:
 - After you create the Kubernetes manifest files, make sure that the **spec.mastersSchedulable** parameter is set to **true** in **cluster-scheduler-02-config.yml** file. You can locate this file in **<installation_directory>/manifests**. For more information, see "Creating the Kubernetes manifest and Ignition config files" in "Installing a cluster on user-provisioned infrastructure in GCP by using Deployment Manager templates".
 - Do not create additional worker nodes.

Example **cluster-scheduler-02-config.yml** file for a three-node cluster

```
apiVersion: config.openshift.io/v1
kind: Scheduler
metadata:
  creationTimestamp: null
```

```

name: cluster
spec:
  mastersSchedulable: true
  policy:
    name: ""
  status: {}

```

9.13.2. Next steps

- [Installing a cluster on GCP with customizations](#)
- [Installing a cluster on user-provisioned infrastructure in GCP by using Deployment Manager templates](#)

9.14. INSTALLATION CONFIGURATION PARAMETERS FOR GCP

Before you deploy an OpenShift Container Platform cluster on Google Cloud Platform (GCP), you provide parameters to customize your cluster and the platform that hosts it. When you create the **install-config.yaml** file, you provide values for the required parameters through the command line. You can then modify the **install-config.yaml** file to customize your cluster further.

9.14.1. Available installation configuration parameters for GCP

The following tables specify the required, optional, and GCP-specific installation configuration parameters that you can set as part of the installation process.



NOTE

After installation, you cannot modify these parameters in the **install-config.yaml** file.

9.14.1.1. Required configuration parameters

Required installation configuration parameters are described in the following table:

Table 9.48. Required parameters

Parameter	Description	Values
apiVersion:	The API version for the install-config.yaml content. The current version is v1 . The installation program may also support older API versions.	String

Parameter	Description	Values
<code>baseDomain:</code>	The base domain of your cloud provider. The base domain is used to create routes to your OpenShift Container Platform cluster components. The full DNS name for your cluster is a combination of the baseDomain and metadata.name parameter values that uses the <metadata.name>.<baseDomain> format.	A fully-qualified domain or subdomain name, such as example.com .
<code>metadata:</code>	Kubernetes resource ObjectMeta , from which only the name parameter is consumed.	Object
<code>metadata: name:</code>	The name of the cluster. DNS records for the cluster are all subdomains of {{.metadata.name}} . {{.baseDomain}} .	String of lowercase letters, hyphens (-), and periods (.), such as dev .
<code>platform:</code>	The configuration for the specific platform upon which to perform the installation: aws, baremetal, azure, gcp, ibmcloud, nutanix, openstack, powersv, vsphere , or {} . For additional information about platform . <platform> parameters, consult the table for your specific platform that follows.	Object
<code>pullSecret:</code>	Get a pull secret from Red Hat OpenShift Cluster Manager to authenticate downloading container images for OpenShift Container Platform components from services such as Quay.io.	<pre>{ "auths":{ "cloud.openshift.com":{ "auth":"b3Blb=", "email":"you@example.com" }, "quay.io":{ "auth":"b3Blb=", "email":"you@example.com" } } }</pre>

9.14.1.2. Network configuration parameters

You can customize your installation configuration based on the requirements of your existing network infrastructure. For example, you can expand the IP address block for the cluster network or provide different IP address blocks than the defaults.


Only IPv4 addresses are supported.




NOTE

Globalnet is not supported with Red Hat OpenShift Data Foundation disaster recovery solutions. For regional disaster recovery scenarios, ensure that you use a nonoverlapping range of private IP addresses for the cluster and service networks in each cluster.

Table 9.49. Network parameters

Parameter	Description	Values
<code>networking:</code>	The configuration for the cluster network.	Object  NOTE You cannot modify parameters specified by the networking object after installation.
<code>networking: networkType:</code>	The Red Hat OpenShift Networking network plugin to install.	OVNKubernetes. OVNKubernetes is a CNI plugin for Linux networks and hybrid networks that contain both Linux and Windows servers. The default value is OVNKubernetes .
<code>networking: clusterNetwork:</code>	The IP address blocks for pods. The default value is 10.128.0.0/14 with a host prefix of /23 . If you specify multiple IP address blocks, the blocks must not overlap.	An array of objects. For example: <code>networking: clusterNetwork: - cidr: 10.128.0.0/14 hostPrefix: 23</code>
<code>networking: clusterNetwork: cidr:</code>	Required if you use networking.clusterNetwork . An IP address block. An IPv4 network.	An IP address block in Classless Inter-Domain Routing (CIDR) notation. The prefix length for an IPv4 block is between 0 and 32 .

Parameter	Description	Values
<code>networking: clusterNetwork: hostPrefix:</code>	The subnet prefix length to assign to each individual node. For example, if hostPrefix is set to 23 then each node is assigned a /23 subnet out of the given cidr . A hostPrefix value of 23 provides 510 ($2^{(32 - 23)} - 2$) pod IP addresses.	A subnet prefix. The default value is 23 .
<code>networking: serviceNetwork:</code>	The IP address block for services. The default value is 172.30.0.0/16 . The OVN-Kubernetes network plugins supports only a single IP address block for the service network.	An array with an IP address block in CIDR format. For example: <code>networking: serviceNetwork: - 172.30.0.0/16</code>
<code>networking: machineNetwork:</code>	The IP address blocks for machines. If you specify multiple IP address blocks, the blocks must not overlap.	An array of objects. For example: <code>networking: machineNetwork: - cidr: 10.0.0.0/16</code>
<code>networking: machineNetwork: cidr:</code>	Required if you use networking.machineNetwork . An IP address block. The default value is 10.0.0.0/16 for all platforms other than libvirt and IBM Power® Virtual Server. For libvirt, the default value is 192.168.126.0/24 . For IBM Power® Virtual Server, the default value is 192.168.0.0/24 .	An IP network block in CIDR notation. For example, 10.0.0.0/16 .  NOTE Set the networking.machineNetwork to match the CIDR that the preferred NIC resides in.


9.14.1.3. Optional configuration parameters


Optional installation configuration parameters are described in the following table:

Table 9.50. Optional parameters

Parameter	Description	Values
<code>additionalTrustBundle:</code>	A PEM-encoded X.509 certificate bundle that is added to the nodes' trusted certificate store. This trust bundle may also be used when a proxy has been configured.	String


Parameter	Description	Values
capabilities:	Controls the installation of optional core cluster components. You can reduce the footprint of your OpenShift Container Platform cluster by disabling optional components. For more information, see the "Cluster capabilities" page in <i>Installing</i> .	String array
capabilities: baselineCapabilitySet:	Selects an initial set of optional capabilities to enable. Valid values are None , v4.11 , v4.12 and vCurrent . The default value is vCurrent .	String
capabilities: additionalEnabledCapabilities:	Extends the set of optional capabilities beyond what you specify in baselineCapabilitySet . You may specify multiple capabilities in this parameter.	String array
cpuPartitioningMode:	Enables workload partitioning, which isolates OpenShift Container Platform services, cluster management workloads, and infrastructure pods to run on a reserved set of CPUs. Workload partitioning can only be enabled during installation and cannot be disabled after installation. While this field enables workload partitioning, it does not configure workloads to use specific CPUs. For more information, see the <i>Workload partitioning</i> page in the <i>Scalability and Performance</i> section.	None or AllNodes . None is the default value.
compute:	The configuration for the machines that comprise the compute nodes.	Array of MachinePool objects.
compute: architecture:	Determines the instruction set architecture of the machines in the pool. Currently, clusters with varied architectures are not supported. All pools must specify the same architecture. Valid values are amd64 and arm64 .	String

Parameter	Description	Values
<code>compute: hyperthreading:</code>	<p>Whether to enable or disable simultaneous multithreading, or hyperthreading, on compute machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.</p>  <p>IMPORTANT</p> <p>If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.</p>	Enabled or Disabled
<code>compute: name:</code>	Required if you use compute . The name of the machine pool.	worker
<code>compute: platform:</code>	Required if you use compute . Use this parameter to specify the cloud provider to host the worker machines. This parameter value must match the controlPlane.platform parameter value.	aws, azure, gcp, ibmcloud, nutanix, openstack, powervs, vsphere, or {}
<code>compute: replicas:</code>	The number of compute machines, which are also known as worker machines, to provision.	A positive integer greater than or equal to 2 . The default value is 3 .
<code>featureSet:</code>	Enables the cluster for a feature set. A feature set is a collection of OpenShift Container Platform features that are not enabled by default. For more information about enabling a feature set during installation, see "Enabling features using feature gates".	String. The name of the feature set to enable, such as TechPreviewNoUpgrade .
<code>controlPlane:</code>	The configuration for the machines that comprise the control plane.	Array of MachinePool objects.


Parameter	Description	Values
controlPlane: architecture:	Determines the instruction set architecture of the machines in the pool. Currently, clusters with varied architectures are not supported. All pools must specify the same architecture. Valid values are amd64 and arm64 .	String
controlPlane: hyperthreading:	Whether to enable or disable simultaneous multithreading, or hyperthreading , on control plane machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.  IMPORTANT If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.	Enabled or Disabled
controlPlane: name:	Required if you use controlPlane . The name of the machine pool.	master
controlPlane: platform:	Required if you use controlPlane . Use this parameter to specify the cloud provider that hosts the control plane machines. This parameter value must match the compute.platform parameter value.	aws, azure, gcp, ibmcloud, nutanix, openstack, powervs, vsphere, or {}
controlPlane: replicas:	The number of control plane machines to provision.	Supported values are 3 , or 1 when deploying single-node OpenShift.

Parameter	Description	Values
credentialsMode:	The Cloud Credential Operator (CCO) mode. If no mode is specified, the CCO dynamically tries to determine the capabilities of the provided credentials, with a preference for mint mode on the platforms where multiple modes are supported.	Mint, Passthrough, Manual or an empty string (""). ^[1]


Parameter	Description	Values
<p>fips:</p>	<p>Enable or disable FIPS mode. The default is false (disabled). If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.</p> <div data-bbox="486 586 592 1637" style="background-color: black; color: white; padding: 10px; margin: 10px 0;"> <p>IMPORTANT</p> <p>To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see Installing the system in FIPS mode.</p> <p>When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.</p> </div> <div data-bbox="486 1686 592 1883" style="background-color: black; color: white; padding: 10px; margin: 10px 0;"> <p>NOTE</p> <p>If you are using Azure File storage, you cannot enable FIPS mode.</p> </div>	<p>false or true</p>

Parameter	Description	Values
<code>imageContentSources:</code>	Sources and repositories for the release-image content.	Array of objects. Includes a source and, optionally, mirrors , as described in the following rows of this table.
<code>imageContentSources: source:</code>	Required if you use imageContentSources . Specify the repository that users refer to, for example, in image pull specifications.	String
<code>imageContentSources: mirrors:</code>	Specify one or more repositories that may also contain the same images.	Array of strings
<code>publish:</code>	How to publish or expose the user-facing endpoints of your cluster, such as the Kubernetes API, OpenShift routes.	Internal or External . To deploy a private cluster, which cannot be accessed from the internet, set publish to Internal . The default value is External .
<code>sshKey:</code>	The SSH key to authenticate access to your cluster machines.  NOTE For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your ssh-agent process uses.	For example, sshKey: ssh-ed25519 AAAA...

1. Not all CCO modes are supported for all cloud providers. For more information about CCO modes, see the "Managing cloud provider credentials" entry in the *Authentication and authorization* content.

**NOTE**

If you are installing on GCP into a shared virtual private cloud (VPC), **credentialsMode** must be set to **Passthrough** or **Manual**.



IMPORTANT

Setting this parameter to **Manual** enables alternatives to storing administrator-level secrets in the **kube-system** project, which require additional configuration steps. For more information, see "Alternatives to storing administrator-level secrets in the kube-system project".

9.14.1.4. Additional Google Cloud Platform (GCP) configuration parameters

Additional GCP configuration parameters are described in the following table:

Table 9.51. Additional GCP parameters

Parameter	Description	Values
controlPlane: platform: gcp: osImage: project:	Optional. By default, the installation program downloads and installs the Red Hat Enterprise Linux CoreOS (RHCOS) image that is used to boot control plane machines. You can override the default behavior by specifying the location of a custom RHCOS image that the installation program is to use for control plane machines only.	String. The name of GCP project where the image is located.

Parameter	Description	Values
<code>controlPlane: platform: gcp: osimage: name:</code>	The name of the custom RHCOS image that the installation program is to use to boot control plane machines. If you use controlPlane.platform.gcp.osimage.project , this field is required.	String. The name of the RHCOS image.
<code>compute: platform: gcp: osimage: project:</code>	Optional. By default, the installation program downloads and installs the RHCOS image that is used to boot compute machines. You can override the default behavior by specifying the location of a custom RHCOS image that the installation program is to use for compute machines only.	String. The name of GCP project where the image is located.

Parameter	Description	Values
<code>compute:platform:gcp:osimage:name:</code>	The name of the custom RHCOS image that the installation program is to use to boot compute machines. If you use compute.platform.gcp.osimage.project , this field is required.	String. The name of the RHCOS image.
<code>platform:gcp:network:</code>	The name of the existing Virtual Private Cloud (VPC) where you want to deploy your cluster. If you want to deploy your cluster into a shared VPC, you must set platform.gcp.networkProjectID with the name of the GCP project that contains the shared VPC.	String.
<code>platform:gcp:networkProjectID:</code>	Optional. The name of the GCP project that contains the shared VPC where you want to deploy your cluster.	String.

Parameter	Description	Values
<code>platform:gcp:projectId</code>	The name of the GCP project where the installation program installs the cluster.	String.
<code>platform:gcp:region</code>	The name of the GCP region that hosts your cluster.	Any valid region name, such as us-central1 .
<code>platform:gcp:controlPlaneSubnet</code>	The name of the existing subnet where you want to deploy your control plane machines.	The subnet name.

Parameter	Description	Values
platform: gcp: compute Subnet:	The name of the existing subnet where you want to deploy your compute machines.	The subnet name.
platform: gcp: default Machine Platform: zones:	The availability zones where the installation program creates machines.	<p>A list of valid GCP availability zones, such as us-central1-a, in a YAML sequence.</p> <div style="display: flex; align-items: flex-start;"> <div style="background-color: black; width: 20px; height: 100%; margin-right: 10px;"></div> <div> <p>IMPORTANT</p> <p>When running your cluster on GCP 64-bit ARM infrastructures, ensure that you use a zone where Ampere Altra Arm CPU's are available. You can find which zones are compatible with 64-bit ARM processors in the "GCP availability zones" link.</p> </div> </div>

Parameter	Description	Values
platform: defaultMachinePlatform: osDisk: diskSizeGB:	The size of the disk in gigabytes (GB).	Any size between 16 GB and 65536 GB.

Parameter	Description	Values
<pre>platform: gcp: defaultMachinePlatform: osDisk: diskType:</pre>	<p>The GCP disk type.</p>	<p>The default disk type for all machines. Control plane nodes must use the pd-ssd disk type. Compute nodes can use the pd-ssd, pd-balanced, or pd-standard disk types.</p>
<pre>platform: gcp: defaultMachinePlatform: osImage: project:</pre>	<p>Optional. By default, the installation program downloads and installs the RHCOS image that is used to boot control plane and compute machines. You can override the default behavior by specifying the location of a custom RHCOS image that the installation program is to use for both types of machines.</p>	<p>String. The name of GCP project where the image is located.</p>

Parameter	Description	Values
<pre>platform: gcp: defaultMachinePlatform: osImageName:</pre>	<p>The name of the custom RHCOS image that the installation program is to use to boot control plane and compute machines. If you use platform.gcp.defaultMachinePlatform.osImage.project, this field is required.</p>	<p>String. The name of the RHCOS image.</p>
<pre>platform: gcp: defaultMachinePlatform: tags:</pre>	<p>Optional. Additional network tags to add to the control plane and compute machines.</p>	<p>One or more strings, for example network-tag1.</p>

Parameter	Description	Values
<pre>platform: gcp: default Machine Platform: type:</pre>	<p>The GCP machine type for control plane and compute machines.</p>	<p>The GCP machine type, for example n1-standard-4.</p>

Parameter	Description	Values
<code>platform: gcp default MachinePlatform: osDisk: encryptionKey: kmsKey: name:</code>	The name of the customer managed encryption key to be used for machine disk encryption.	The encryption key name.

Parameter	Description	Values
<pre>platform: gcp: default MachinePlatform: osDisk: encryptionKey: kmsKeyRing:</pre>	<p>The name of the Key Management Service (KMS) key ring to which the KMS key belongs.</p>	<p>The KMS key ring name.</p>

Parameter	Description	Values
<code>platform: gcp default MachinePlatform: osDisk: encryptionKey: kmsKey: location:</code>	The GCP location in which the KMS key ring exists.	The GCP location.

Parameter	Description	Values
platform: gcp: default Machine Platform: os Disk: encryption Key: kmsKey: projectId:	The ID of the project in which the KMS key ring exists. This value defaults to the value of the platform.gcp.projectID parameter if it is not set.	The GCP project ID.

Parameter	Description	Values
<code>platform: gcp: defaultMachinePlatform: osDisk: encryptionKey: kmsKeyServiceAccount:</code>	<p>The GCP service account used for the encryption request for control plane and compute machines. If absent, the Compute Engine default service account is used. For more information about GCP service accounts, see Google's documentation on service accounts.</p>	<p>The GCP service account email, for example <service_account_name>@<project_id>.iam.gserviceaccount.com.</p>

Parameter	Description	Values
<pre>platform: gcp: defaultMachinePlatform: secureBoot:</pre>	<p>Whether to enable Shielded VM secure boot for all machines in the cluster. Shielded VMs have additional security protocols such as secure boot, firmware and integrity monitoring, and rootkit protection. For more information on Shielded VMs, see Google's documentation on Shielded VMs.</p>	<p>Enabled or Disabled. The default value is Disabled.</p>
<pre>platform: gcp: defaultMachinePlatform: confidentialCompute:</pre>	<p>Whether to use Confidential VMs for all machines in the cluster. Confidential VMs provide encryption for data during processing. For more information on Confidential computing, see Google's documentation on Confidential computing.</p>	<p>Enabled or Disabled. The default value is Disabled.</p>

Parameter	Description	Values
<code>platform.gcp.defaultMachinePlatform.onHostMaintenance:</code>	Specifies the behavior of all VMs during a host maintenance event, such as a software or hardware update. For Confidential VMs, this parameter must be set to Terminate . Confidential VMs do not support live VM migration.	Terminate or Migrate . The default value is Migrate .

Parameter	Description	Values
control Plane: platform: gcp: os Disk: encryption Key: ksKey: name:	The name of the customer managed encryption key to be used for control plane machine disk encryption.	The encryption key name.

Parameter	Description	Values
controlPlane:platform:gcp:osDisk:encryptionKey:kmsKeyRing:	For control plane machines, the name of the KMS key ring to which the KMS key belongs.	The KMS key ring name.


Parameter	Description	Values
control Plane: platform: gcp: os Disk: encryption Key: kmsKey: location:	For control plane machines, the GCP location in which the key ring exists. For more information about KMS locations, see Google's documentation on Cloud KMS locations .	The GCP location for the key ring.

Parameter	Description	Values
controlPlane: platform: gcp: osDisk: encryptionKey: kmsKey: projectId:	For control plane machines, the ID of the project in which the KMS key ring exists. This value defaults to the VM project ID if not set.	The GCP project ID.

Parameter	Description	Values
<p>control Plane:</p> <p>platform:</p> <p>gcp:</p> <p>os Disk:</p> <p>encryption Key:</p> <p>km sKey ySe rvic eAc cou nt:</p>	<p>The GCP service account used for the encryption request for control plane machines. If absent, the Compute Engine default service account is used. For more information about GCP service accounts, see Google's documentation on service accounts.</p>	<p>The GCP service account email, for example <service_account_name>@<project_id>.iam.gserviceaccount.com.</p>

Parameter	Description	Values
controlPlane: platform: gcp: osDisk: diskSizeGB:	The size of the disk in gigabytes (GB). This value applies to control plane machines.	Any integer between 16 and 65536.
controlPlane: platform: gcp: osDisk: diskType:	The GCP disk type for control plane machines.	Control plane machines must use the pd-ssd disk type, which is the default.

Parameter	Description	Values
<code>controlPlane: platform: gcp: tags:</code>	<p>Optional. Additional network tags to add to the control plane machines. If set, this parameter overrides the <code>platform.gcp.defaultMachinePlatform.tags</code> parameter for control plane machines.</p>	<p>One or more strings, for example <code>control-plane-tag1</code>.</p>
<code>controlPlane: platform: gcp: type:</code>	<p>The GCP machine type for control plane machines. If set, this parameter overrides the <code>platform.gcp.defaultMachinePlatform.type</code> parameter.</p>	<p>The GCP machine type, for example <code>n1-standard-4</code>.</p>

Parameter	Description	Values
controlPlane: platform: gcp: zones:	The availability zones where the installation program creates control plane machines.	<p>A list of valid GCP availability zones, such as us-central1-a, in a YAML sequence.</p>  <p>IMPORTANT</p> <p>When running your cluster on GCP 64-bit ARM infrastructure, ensure that you use a zone where Ampere Altra Arm CPU's are available. You can find which zones are compatible with 64-bit ARM processors in the "GCP availability zones" link.</p>
controlPlane: platform: gcp: secureBoot:	Whether to enable Shielded VM secure boot for control plane machines. Shielded VMs have additional security protocols such as secure boot, firmware and integrity monitoring, and rootkit protection. For more information on Shielded VMs, see Google's documentation on Shielded VMs .	Enabled or Disabled . The default value is Disabled .

Parameter	Description	Values
<pre>controlPlane: platform: gcp: confidentialCompute:</pre>	<p>Whether to enable Confidential VMs for control plane machines. Confidential VMs provide encryption for data while it is being processed. For more information on Confidential VMs, see Google's documentation on Confidential Computing.</p>	<p>Enabled or Disabled. The default value is Disabled.</p>
<pre>controlPlane: platform: gcp: onHostMaintenance:</pre>	<p>Specifies the behavior of control plane VMs during a host maintenance event, such as a software or hardware update. For Confidential VMs, this parameter must be set to Terminate. Confidential VMs do not support live VM migration.</p>	<p>Terminate or Migrate. The default value is Migrate.</p>

Parameter	Description	Values
compute:platform:gcp:osDisk:encryptionKey:keyname:	The name of the customer managed encryption key to be used for compute machine disk encryption.	The encryption key name.


Parameter	Description	Values
compute:platform:gcp:osDisk:encryptionKey:keyRing:	For compute machines, the name of the KMS key ring to which the KMS key belongs.	The KMS key ring name.

Parameter	Description	Values
computePlatform.gcpDiskEncryptionKey:keyLocation	For compute machines, the GCP location in which the key ring exists. For more information about KMS locations, see Google's documentation on Cloud KMS locations .	The GCP location for the key ring.

Parameter	Description	Values
compute: platform: gcp: os Disk: encryption Key: kmsKey: projectId:	For compute machines, the ID of the project in which the KMS key ring exists. This value defaults to the VM project ID if not set.	The GCP project ID.

Parameter	Description	Values
<code>compute:platform:gcp:osDisk:encryptionKey:skeserviceAccount:</code>	The GCP service account used for the encryption request for compute machines. If this value is not set, the Compute Engine default service account is used. For more information about GCP service accounts, see Google's documentation on service accounts .	The GCP service account email, for example <service_account_name>@<project_id>.iam.gserviceaccount.com .
<code>compute:platform:gcp:osDisk:diskSizeGB:</code>	The size of the disk in gigabytes (GB). This value applies to compute machines.	Any integer between 16 and 65536.

Parameter	Description	Values
<code>compute:platform:gcp:osDisk:diskType:</code>	The GCP disk type for compute machines.	pd-ssd , pd-standard , or pd-balanced . The default is pd-ssd .
<code>compute:platform:gcp:tags:</code>	Optional. Additional network tags to add to the compute machines. If set, this parameter overrides the platform.gcp.defaultMachinePlatform.tags parameter for compute machines.	One or more strings, for example compute-network-tag1 .
<code>compute:platform:gcp:type:</code>	The GCP machine type for compute machines. If set, this parameter overrides the platform.gcp.defaultMachinePlatform.type parameter.	The GCP machine type, for example n1-standard-4 .

Parameter	Description	Values
<pre>compute: platform: gcp: zones:</pre>	<p>The availability zones where the installation program creates compute machines.</p>	<p>A list of valid GCP availability zones, such as us-central1-a, in a YAML sequence.</p> <div style="display: flex; align-items: center;">  <div style="flex: 1;"> <p>IMPORTANT</p> <p>When running your cluster on GCP 64-bit ARM infrastructures, ensure that you use a zone where Ampere Altra Arm CPU's are available. You can find which zones are compatible with 64-bit ARM processors in the "GCP availability zones" link.</p> </div> </div>
<pre>compute: platform: gcp: secureBoot:</pre>	<p>Whether to enable Shielded VM secure boot for compute machines. Shielded VMs have additional security protocols such as secure boot, firmware and integrity monitoring, and rootkit protection. For more information on Shielded VMs, see Google's documentation on Shielded VMs.</p>	<p>Enabled or Disabled. The default value is Disabled.</p>

Parameter	Description	Values
compute:platform:gcpcnfidentialICompute:	Whether to enable Confidential VMs for compute machines. Confidential VMs provide encryption for data while it is being processed. For more information on Confidential VMs, see Google's documentation on Confidential Computing .	Enabled or Disabled . The default value is Disabled .
compute:platform:gcponHostMaintenance:	Specifies the behavior of compute VMs during a host maintenance event, such as a software or hardware update. For Confidential VMs, this parameter must be set to Terminate . Confidential VMs do not support live VM migration.	Terminate or Migrate . The default value is Migrate .

9.15. UNINSTALLING A CLUSTER ON GCP

You can remove a cluster that you deployed to Google Cloud Platform (GCP).

9.15.1. Removing a cluster that uses installer-provisioned infrastructure

You can remove a cluster that uses installer-provisioned infrastructure from your cloud.



NOTE

After uninstallation, check your cloud provider for any resources not removed properly, especially with User Provisioned Infrastructure (UPI) clusters. There might be resources that the installer did not create or that the installer is unable to access. For example, some Google Cloud resources require [IAM permissions](#) in shared VPC host projects, or there might be unused [health checks that must be deleted](#).

Prerequisites

- You have a copy of the installation program that you used to deploy the cluster.
- You have the files that the installation program generated when you created your cluster.

Procedure

1. From the directory that contains the installation program on the computer that you used to install the cluster, run the following command:

```
$ ./openshift-install destroy cluster \
--dir <installation_directory> --log-level info 1 2
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.
- 2** To view different details, specify **warn**, **debug**, or **error** instead of **info**.



NOTE

You must specify the directory that contains the cluster definition files for your cluster. The installation program requires the **metadata.json** file in this directory to delete the cluster.

2. Optional: Delete the **<installation_directory>** directory and the OpenShift Container Platform installation program.

9.15.2. Deleting Google Cloud Platform resources with the Cloud Credential Operator utility

After uninstalling an OpenShift Container Platform cluster that uses short-term credentials managed outside the cluster, you can use the CCO utility (**ccoctl**) to remove the Google Cloud Platform (GCP) resources that **ccoctl** created during installation.

Prerequisites

- Extract and prepare the **ccoctl** binary.
- Uninstall an OpenShift Container Platform cluster on GCP that uses short-term credentials.

Procedure

1. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

2. Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included \1
  --to=<path_to_directory_for_credentials_requests> \2
```

- 1 The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2 Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

3. Delete the GCP resources that **ccoctl** created by running the following command:

```
$ ccoctl gcp delete \
  --name=<name> \1
  --project=<gcp_project_id> \2
  --credentials-requests-dir=<path_to_credentials_requests_directory> \
  --force-delete-custom-roles \3
```

- 1 **<name>** matches the name that was originally used to create and tag the cloud resources.
- 2 **<gcp_project_id>** is the GCP project ID in which to delete cloud resources.
- 3 Optional: This parameter deletes the custom roles that the **ccoctl** utility creates during installation. GCP does not permanently delete custom roles immediately. For more information, see GCP documentation about [deleting a custom role](#).

Verification

- To verify that the resources are deleted, query GCP. For more information, refer to GCP documentation.

CHAPTER 10. INSTALLING ON IBM CLOUD

10.1. PREPARING TO INSTALL ON IBM CLOUD

The installation workflows documented in this section are for IBM Cloud® infrastructure environments. IBM Cloud® classic is not supported at this time. For more information about the difference between classic and VPC infrastructures, see the IBM® [documentation](#).

10.1.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).

10.1.2. Requirements for installing OpenShift Container Platform on IBM Cloud

Before installing OpenShift Container Platform on IBM Cloud®, you must create a service account and configure an IBM Cloud® account. See [Configuring an IBM Cloud® account](#) for details about creating an account, enabling API services, configuring DNS, IBM Cloud® account limits, and supported IBM Cloud® regions.

You must manually manage your cloud credentials when installing a cluster to IBM Cloud®. Do this by configuring the Cloud Credential Operator (CCO) for manual mode before you install the cluster. For more information, see [Configuring IAM for IBM Cloud®](#).

10.1.3. Choosing a method to install OpenShift Container Platform on IBM Cloud

You can install OpenShift Container Platform on IBM Cloud® using installer-provisioned infrastructure. This process involves using an installation program to provision the underlying infrastructure for your cluster. Installing OpenShift Container Platform on IBM Cloud® using user-provisioned infrastructure is not supported at this time.

See [Installation process](#) for more information about installer-provisioned installation processes.

10.1.3.1. Installing a cluster on installer-provisioned infrastructure

You can install a cluster on IBM Cloud® infrastructure that is provisioned by the OpenShift Container Platform installation program by using one of the following methods:

- **Installing a customized cluster on IBM Cloud®** You can install a customized cluster on IBM Cloud® infrastructure that the installation program provisions. The installation program allows for some customization to be applied at the installation stage. Many other customization options are available [post-installation](#).
- **Installing a cluster on IBM Cloud® with network customizations** You can customize your OpenShift Container Platform network configuration during installation, so that your cluster can coexist with your existing IP address allocations and adhere to your network requirements.
- **Installing a cluster on IBM Cloud® into an existing VPC** You can install OpenShift Container Platform on an existing IBM Cloud®. You can use this installation method if you have constraints set by the guidelines of your company, such as limits when creating new accounts or infrastructure.

- **Installing a private cluster on an existing VPC** You can install a private cluster on an existing Virtual Private Cloud (VPC). You can use this method to deploy OpenShift Container Platform on an internal network that is not visible to the internet.
- **Installing a cluster on IBM Cloud VPC in a restricted network** You can install OpenShift Container Platform on IBM Cloud VPC on installer-provisioned infrastructure by using an internal mirror of the installation release content. You can use this method to install a cluster that does not require an active internet connection to obtain the software components.

10.1.4. Next steps

- [Configuring an IBM Cloud® account](#)

10.2. CONFIGURING AN IBM CLOUD ACCOUNT

Before you can install OpenShift Container Platform, you must configure an IBM Cloud® account.

10.2.1. Prerequisites

- You have an IBM Cloud® account with a subscription. You cannot install OpenShift Container Platform on a free or trial IBM Cloud® account.

10.2.2. Quotas and limits on IBM Cloud

The OpenShift Container Platform cluster uses a number of IBM Cloud® components, and the default quotas and limits affect your ability to install OpenShift Container Platform clusters. If you use certain cluster configurations, deploy your cluster in certain regions, or run multiple clusters from your account, you might need to request additional resources for your IBM Cloud® account.

For a comprehensive list of the default IBM Cloud® quotas and service limits, see IBM Cloud®'s documentation for [Quotas and service limits](#).

Virtual Private Cloud (VPC)

Each OpenShift Container Platform cluster creates its own VPC. The default quota of VPCs per region is 10 and will allow 10 clusters. To have more than 10 clusters in a single region, you must increase this quota.

Application load balancer

By default, each cluster creates three application load balancers (ALBs):

- Internal load balancer for the master API server
- External load balancer for the master API server
- Load balancer for the router

You can create additional **LoadBalancer** service objects to create additional ALBs. The default quota of VPC ALBs are 50 per region. To have more than 50 ALBs, you must increase this quota.

VPC ALBs are supported. Classic ALBs are not supported for IBM Cloud®.

Floating IP address

By default, the installation program distributes control plane and compute machines across all availability zones within a region to provision the cluster in a highly available configuration. In each availability zone, a public gateway is created and requires a separate floating IP address.

The default quota for a floating IP address is 20 addresses per availability zone. The default cluster configuration yields three floating IP addresses:

- Two floating IP addresses in the **us-east-1** primary zone. The IP address associated with the bootstrap node is removed after installation.
- One floating IP address in the **us-east-2** secondary zone.
- One floating IP address in the **us-east-3** secondary zone.

IBM Cloud® can support up to 19 clusters per region in an account. If you plan to have more than 19 default clusters, you must increase this quota.

Virtual Server Instances (VSI)

By default, a cluster creates VSIs using **bx2-4x16** profiles, which includes the following resources by default:

- 4 vCPUs
- 16 GB RAM

The following nodes are created:

- One **bx2-4x16** bootstrap machine, which is removed after the installation is complete
- Three **bx2-4x16** control plane nodes
- Three **bx2-4x16** compute nodes

For more information, see IBM Cloud®'s documentation on [supported profiles](#).

Table 10.1. VSI component quotas and limits

VSI component	Default IBM Cloud® quota	Default cluster configuration	Maximum number of clusters
vCPU	200 vCPUs per region	28 vCPUs, or 24 vCPUs after bootstrap removal	8 per region
RAM	1600 GB per region	112 GB, or 96 GB after bootstrap removal	16 per region
Storage	18 TB per region	1050 GB, or 900 GB after bootstrap removal	19 per region

If you plan to exceed the resources stated in the table, you must increase your IBM Cloud® account quota.

Block Storage Volumes

For each VPC machine, a block storage device is attached for its boot volume. The default cluster configuration creates seven VPC machines, resulting in seven block storage volumes. Additional Kubernetes persistent volume claims (PVCs) of the IBM Cloud® storage class create additional block storage volumes. The default quota of VPC block storage volumes are 300 per region. To have more than 300 volumes, you must increase this quota.


10.2.3. Configuring DNS resolution

How you configure DNS resolution depends on the type of OpenShift Container Platform cluster you are installing:

- If you are installing a public cluster, you use IBM Cloud Internet Services (CIS).
- If you are installing a private cluster, you use IBM Cloud® DNS Services (DNS Services)

10.2.3.1. Using IBM Cloud Internet Services for DNS resolution

The installation program uses IBM Cloud® Internet Services (CIS) to configure cluster DNS resolution and provide name lookup for a public cluster.



NOTE

This offering does not support IPv6, so dual stack or IPv6 environments are not possible.

You must create a domain zone in CIS in the same account as your cluster. You must also ensure the zone is authoritative for the domain. You can do this using a root domain or subdomain.

Prerequisites

- You have installed the [IBM Cloud® CLI](#).
- You have an existing domain and registrar. For more information, see the [IBM® documentation](#).

Procedure

1. Create a CIS instance to use with your cluster:

a. Install the CIS plugin:

```
$ ibmcloud plugin install cis
```

b. Create the CIS instance:

```
$ ibmcloud cis instance-create <instance_name> standard 1
```

1 At a minimum, a **Standard** plan is required for CIS to manage the cluster subdomain and its DNS records.

2. Connect an existing domain to your CIS instance:

a. Set the context instance for CIS:

```
$ ibmcloud cis instance-set <instance_name> 1
```

1 The instance cloud resource name.

b. Add the domain for CIS:

```
$ ibmcloud cis domain-add <domain_name> 1
```

- 1** The fully qualified domain name. You can use either the root domain or subdomain value as the domain name, depending on which you plan to configure.



NOTE

A root domain uses the form **openshiftcorp.com**. A subdomain uses the form **clusters.openshiftcorp.com**.

3. Open the [CIS web console](#), navigate to the **Overview** page, and note your CIS name servers. These name servers will be used in the next step.
4. Configure the name servers for your domains or subdomains at the domain's registrar or DNS provider. For more information, see the IBM Cloud® [documentation](#).

10.2.3.2. Using IBM Cloud DNS Services for DNS resolution

The installation program uses IBM Cloud® DNS Services to configure cluster DNS resolution and provide name lookup for a private cluster.

You configure DNS resolution by creating a DNS services instance for the cluster, and then adding a DNS zone to the DNS Services instance. Ensure that the zone is authoritative for the domain. You can do this using a root domain or subdomain.



NOTE

IBM Cloud® does not support IPv6, so dual stack or IPv6 environments are not possible.

Prerequisites

- You have installed the [IBM Cloud® CLI](#).
- You have an existing domain and registrar. For more information, see the IBM® [documentation](#).

Procedure

1. Create a DNS Services instance to use with your cluster:
 - a. Install the DNS Services plugin by running the following command:

```
$ ibmcloud plugin install cloud-dns-services
```

- b. Create the DNS Services instance by running the following command:

```
$ ibmcloud dns instance-create <instance-name> standard-dns 1
```

- 1** At a minimum, a **Standard** plan is required for DNS Services to manage the cluster subdomain and its DNS records.

2. Create a DNS zone for the DNS Services instance:

- a. Set the target operating DNS Services instance by running the following command:

```
$ ibmcloud dns instance-target <instance-name>
```

- b. Add the DNS zone to the DNS Services instance by running the following command:

```
$ ibmcloud dns zone-create <zone-name> 1
```

- 1** The fully qualified zone name. You can use either the root domain or subdomain value as the zone name, depending on which you plan to configure. A root domain uses the form **openshiftcorp.com**. A subdomain uses the form **clusters.openshiftcorp.com**.

3. Record the name of the DNS zone you have created. As part of the installation process, you must update the **install-config.yaml** file before deploying the cluster. Use the name of the DNS zone as the value for the **baseDomain** parameter.



NOTE

You do not have to manage permitted networks or configure an "A" DNS resource record. As required, the installation program configures these resources automatically.

10.2.4. IBM Cloud IAM Policies and API Key

To install OpenShift Container Platform into your IBM Cloud® account, the installation program requires an IAM API key, which provides authentication and authorization to access IBM Cloud® service APIs. You can use an existing IAM API key that contains the required policies or create a new one.

For an IBM Cloud® IAM overview, see the IBM Cloud® [documentation](#).

10.2.4.1. Required access policies

You must assign the required access policies to your IBM Cloud® account.

Table 10.2. Required access policies

Service type	Service	Access policy scope	Platform access	Service access
Account management	IAM Identity Service	All resources or a subset of resources ^[1]	Editor, Operator, Viewer, Administrator	Service ID creator
Account management ^[2]	Identity and Access Management	All resources	Editor, Operator, Viewer, Administrator	

Service type	Service	Access policy scope	Platform access	Service access
Account management	Resource group only	All resource groups in the account	Administrator	
IAM services	Cloud Object Storage	All resources or a subset of resources ^[1]	Editor, Operator, Viewer, Administrator	Reader, Writer, Manager, Content Reader, Object Reader, Object Writer
IAM services	Internet Services	All resources or a subset of resources ^[1]	Editor, Operator, Viewer, Administrator	Reader, Writer, Manager
IAM services	DNS Services	All resources or a subset of resources ^[1]	Editor, Operator, Viewer, Administrator	Reader, Writer, Manager
IAM services	VPC Infrastructure Services	All resources or a subset of resources ^[1]	Editor, Operator, Viewer, Administrator	Reader, Writer, Manager

1. The policy access scope should be set based on how granular you want to assign access. The scope can be set to **All resources** or **Resources based on selected attributes**
2. Optional: This access policy is only required if you want the installation program to create a resource group. For more information about resource groups, see the IBM® [documentation](#).

10.2.4.2. Access policy assignment

In IBM Cloud® IAM, access policies can be attached to different subjects:

- Access group (Recommended)
- Service ID
- User

The recommended method is to define IAM access policies in an [access group](#). This helps organize all the access required for OpenShift Container Platform and enables you to onboard users and service IDs to this group. You can also assign access to [users and service IDs](#) directly, if desired.

10.2.4.3. Creating an API key

You must create a user API key or a service ID API key for your IBM Cloud® account.

Prerequisites

- You have assigned the required access policies to your IBM Cloud® account.

- You have attached your IAM access policies to an access group, or other appropriate resource.

Procedure

- Create an API key, depending on how you defined your IAM access policies. For example, if you assigned your access policies to a user, you must create a [user API key](#). If you assigned your access policies to a service ID, you must create a [service ID API key](#). If your access policies are assigned to an access group, you can use either API key type. For more information on IBM Cloud® API keys, see [Understanding API keys](#).

10.2.5. Supported IBM Cloud regions

You can deploy an OpenShift Container Platform cluster to the following regions:

- **au-syd** (Sydney, Australia)
- **br-sao** (Sao Paulo, Brazil)
- **ca-tor** (Toronto, Canada)
- **eu-de** (Frankfurt, Germany)
- **eu-gb** (London, United Kingdom)
- **eu-es** (Madrid, Spain)
- **jp-osa** (Osaka, Japan)
- **jp-tok** (Tokyo, Japan)
- **us-east** (Washington DC, United States)
- **us-south** (Dallas, United States)



NOTE

Deploying your cluster in the **eu-es** (Madrid, Spain) region is not supported for OpenShift Container Platform 4.14.6 and earlier versions.

10.2.6. Next steps

- [Configuring IAM for IBM Cloud®](#)

10.3. CONFIGURING IAM FOR IBM CLOUD

In environments where the cloud identity and access management (IAM) APIs are not reachable, you must put the Cloud Credential Operator (CCO) into manual mode before you install the cluster.

10.3.1. Alternatives to storing administrator-level secrets in the kube-system project

The Cloud Credential Operator (CCO) manages cloud provider credentials as Kubernetes custom resource definitions (CRDs). You can configure the CCO to suit the security requirements of your organization by setting different values for the **credentialsMode** parameter in the **install-config.yaml** file.

Storing an administrator-level credential secret in the cluster **kube-system** project is not supported for IBM Cloud®; therefore, you must set the **credentialsMode** parameter for the CCO to **Manual** when installing OpenShift Container Platform and manage your cloud credentials manually.

Using manual mode allows each cluster component to have only the permissions it requires, without storing an administrator-level credential in the cluster. You can also use this mode if your environment does not have connectivity to the cloud provider public IAM endpoint. However, you must manually reconcile permissions with new release images for every upgrade. You must also manually supply credentials for every component that requests them.

Additional resources

- [About the Cloud Credential Operator](#)

10.3.2. Configuring the Cloud Credential Operator utility

To create and manage cloud credentials from outside of the cluster when the Cloud Credential Operator (CCO) is operating in manual mode, extract and prepare the CCO utility (**ccoctl**) binary.



NOTE

The **ccoctl** utility is a Linux binary that must run in a Linux environment.

Prerequisites

- You have access to an OpenShift Container Platform account with cluster administrator access.
- You have installed the OpenShift CLI (**oc**).

Procedure

1. Set a variable for the OpenShift Container Platform release image by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

2. Obtain the CCO container image from the OpenShift Container Platform release image by running the following command:

```
$ CCO_IMAGE=$(oc adm release info --image-for='cloud-credential-operator'
$RELEASE_IMAGE -a ~/.pull-secret)
```



NOTE

Ensure that the architecture of the **\$RELEASE_IMAGE** matches the architecture of the environment in which you will use the **ccoctl** tool.

3. Extract the **ccoctl** binary from the CCO container image within the OpenShift Container Platform release image by running the following command:

```
$ oc image extract $CCO_IMAGE \
--file="/usr/bin/ccoctl.<rhel_version>" \
-a ~/.pull-secret
```

- 1 For **<rhel_version>**, specify the value that corresponds to the version of Red Hat Enterprise Linux (RHEL) that the host uses. If no value is specified, **ccoctl.rhel8** is used by default. The following values are valid:
- **rhel8**: Specify this value for hosts that use RHEL 8.
 - **rhel9**: Specify this value for hosts that use RHEL 9.

4. Change the permissions to make **ccoctl** executable by running the following command:

```
$ chmod 775 ccoctl.<rhel_version>
```

Verification

- To verify that **ccoctl** is ready to use, display the help file by running the following command:

```
$ ccoctl --help
```

Output of **ccoctl --help**

```
OpenShift credentials provisioning tool
```

```
Usage:
```

```
ccoctl [command]
```

```
Available Commands:
```

```
aws      Manage credentials objects for AWS cloud
azure    Manage credentials objects for Azure
gcp      Manage credentials objects for Google cloud
help     Help about any command
ibmcloud Manage credentials objects for IBM Cloud
nutanix  Manage credentials objects for Nutanix
```

```
Flags:
```

```
-h, --help  help for ccoctl
```

```
Use "ccoctl [command] --help" for more information about a command.
```

Additional resources

- [Rotating API keys for IBM Cloud®](#)

10.3.3. Next steps

- [Installing a cluster on IBM Cloud® with customizations](#)

10.3.4. Additional resources

- [Preparing to update a cluster with manually maintained credentials](#)

10.4. USER-MANAGED ENCRYPTION FOR IBM CLOUD

By default, provider-managed encryption is used to secure the following when you deploy an OpenShift Container Platform cluster:

- The root (boot) volume of control plane and compute machines
- Persistent volumes (data volumes) that are provisioned after the cluster is deployed

You can override the default behavior by specifying an IBM® Key Protect for IBM Cloud® (Key Protect) root key as part of the installation process.

When you bring our own root key, you modify the installation configuration file (**install-config.yaml**) to specify the Cloud Resource Name (CRN) of the root key by using the **encryptionKey** parameter.

You can specify that:

- The same root key be used for all cluster machines. You do so by specifying the key as part of the cluster's default machine configuration. When specified as part of the default machine configuration, all managed storage classes are updated with this key. As such, data volumes that are provisioned after the installation are also encrypted using this key.
- Separate root keys be used for the control plane and compute machine pools.

For more information about the **encryptionKey** parameter, see [Additional IBM Cloud configuration parameters](#).



NOTE

Make sure you have integrated Key Protect with your IBM Cloud Block Storage service. For more information, see the Key Protect [documentation](#).

10.4.1. Next steps

Install an OpenShift Container Platform cluster:

- [Installing a cluster on IBM Cloud with customizations](#)
- [Installing a cluster on IBM Cloud with network customizations](#)
- [Installing a cluster on IBM Cloud into an existing VPC](#)
- [Installing a private cluster on IBM Cloud](#)

10.5. INSTALLING A CLUSTER ON IBM CLOUD WITH CUSTOMIZATIONS

In OpenShift Container Platform version 4.16, you can install a customized cluster on infrastructure that the installation program provisions on IBM Cloud®. To customize the installation, you modify parameters in the **install-config.yaml** file before you install the cluster.

10.5.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.


- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [configured an IBM Cloud® account](#) to host the cluster.
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.
- You configured the **ccoctl** utility before you installed the cluster. For more information, see [Configuring IAM for IBM Cloud®](#).

10.5.2. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.



IMPORTANT


If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

10.5.3. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The **`./openshift-install gather`** command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.

**NOTE**

You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.

**NOTE**

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.

**NOTE**

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```


**NOTE**

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

10.5.4. Obtaining the installation program


Before you install OpenShift Container Platform, download the installation file on the host you are using for installation.

Prerequisites


- You have a computer that runs Linux or macOS, with 500 MB of local disk space.

Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select your infrastructure provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.

**IMPORTANT**

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.



IMPORTANT

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.

4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#) . This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

10.5.5. Exporting the API key

You must set the API key you created as a global variable; the installation program ingests the variable during startup to set the API key.

Prerequisites

- You have created either a user API key or service ID API key for your IBM Cloud® account.

Procedure

- Export your API key for your account as a global variable:

```
$ export IC_API_KEY=<api_key>
```



IMPORTANT

You must set the variable name exactly as specified; the installation program expects the variable name to be present during startup.

10.5.6. Creating the installation configuration file

You can customize the OpenShift Container Platform cluster you install on IBM Cloud®.

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create the **install-config.yaml** file.
 - a. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install create install-config --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the directory name to store the files that the installation program creates.

When specifying the directory:

- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
 - Use an empty directory. Some installation assets, such as bootstrap X.509 certificates, have short expiration intervals, therefore you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.
- b. At the prompts, provide the configuration details for your cloud:
- i. Optional: Select an SSH key to use to access your cluster machines.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- ii. Select **ibmcloud** as the platform to target.
 - iii. Select the region to deploy the cluster to.
 - iv. Select the base domain to deploy the cluster to. The base domain corresponds to the public DNS zone that you created for your cluster.
 - v. Enter a descriptive name for your cluster.
2. Modify the **install-config.yaml** file. You can find more information about the available parameters in the "Installation configuration parameters" section.
 3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

Additional resources

- [Installation configuration parameters for IBM Cloud®](#)

10.5.6.1. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 10.3. Minimum resource requirements

Machine	Operating System	vCPU	Virtual RAM	Storage	Input/Output Per Second (IOPS)
Bootstrap	RHCOS	4	16 GB	100 GB	300
Control plane	RHCOS	4	16 GB	100 GB	300
Compute	RHCOS	2	8 GB	100 GB	300

**NOTE**

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

10.5.6.2. Tested instance types for IBM Cloud

The following IBM Cloud® instance types have been tested with OpenShift Container Platform.

Example 10.1. Machine series

- **c4.***
- **c5.***
- **c5a.***
- **i3.***
- **m4.***
- **m5.***
- **m5a.***

- m6a.*
- m6i.*
- r4.*
- r5.*
- r5a.*
- r6i.*
- t3.*
- t3a.*

10.5.6.3. Sample customized install-config.yaml file for IBM Cloud

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.



IMPORTANT

This sample YAML file is provided for reference only. You must obtain your **install-config.yaml** file by using the installation program and then modify it.

```

apiVersion: v1
baseDomain: example.com 1
controlPlane: 2 3
  hyperthreading: Enabled 4
  name: master
  platform:
    ibmcloud: {}
  replicas: 3
compute: 5 6
- hyperthreading: Enabled 7
  name: worker
  platform:
    ibmcloud: {}
  replicas: 3
metadata:
  name: test-cluster 8
networking:
  clusterNetwork:
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  machineNetwork:
  - cidr: 10.0.0.0/16
  networkType: OVNKubernetes 9
  serviceNetwork:
  - 172.30.0.0/16
platform:
  ibmcloud:

```

```

region: us-south 10
credentialsMode: Manual
publish: External
pullSecret: '{"auths": ...}' 11
fips: false 12
sshKey: ssh-ed25519 AAAA... 13

```

- 1 8 10 11** Required. The installation program prompts you for this value.
- 2 5** If you do not provide these parameters and values, the installation program provides the default value.
- 3 6** The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.
- 4 7** Enables or disables simultaneous multithreading, also known as Hyper-Threading. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores. You can disable it by setting the parameter value to **Disabled**. If you disable simultaneous multithreading in some cluster machines, you must disable it in all cluster machines.



IMPORTANT

If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance. Use larger machine types, such as **n1-standard-8**, for your machines if you disable simultaneous multithreading.

- 9** The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 12** Enables or disables FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.



IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 13** Optional: provide the **sshKey** value that you use to access the machines in your cluster.

**NOTE**

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

10.5.6.4. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.

**NOTE**

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
additionalTrustBundle: | 4
  -----BEGIN CERTIFICATE-----
  <MY_TRUSTED_CA_CERT>
  -----END CERTIFICATE-----
additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.

- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with `.` to match subdomains only. For
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.

**NOTE**

The installation program does not support the proxy **readinessEndpoints** field.

**NOTE**

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.

**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

10.5.7. Manually creating IAM

Installing the cluster requires that the Cloud Credential Operator (CCO) operate in manual mode. While the installation program configures the CCO for manual mode, you must specify the identity and access management secrets for your cloud provider.

You can use the Cloud Credential Operator (CCO) utility (**ccoctl**) to create the required IBM Cloud® resources.

Prerequisites

- You have configured the **ccoctl** binary.
- You have an existing **install-config.yaml** file.

Procedure

1. Edit the **install-config.yaml** configuration file so that it contains the **credentialsMode** parameter set to **Manual**.

Example install-config.yaml configuration file

```
apiVersion: v1
baseDomain: cluster1.example.com
credentialsMode: Manual 1
compute:
- architecture: amd64
  hyperthreading: Enabled
```

- 1** This line is added to set the **credentialsMode** parameter to **Manual**.

2. To generate the manifests, run the following command from the directory that contains the installation program:

```
$ openshift-install create manifests --dir <installation_directory>
```

3. From the directory that contains the installation program, set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

4. Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included 1 \
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml 2 \
  --to=<path_to_directory_for_credentials_requests> 3
```

- 1** The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2** Specify the location of the **install-config.yaml** file.
- 3** Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

This command creates a YAML file for each **CredentialsRequest** object.

Sample CredentialsRequest object

```
apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  labels:
```

```

    controller-tools.k8s.io: "1.0"
    name: openshift-image-registry-ibmcos
    namespace: openshift-cloud-credential-operator
spec:
  secretRef:
    name: installer-cloud-credentials
    namespace: openshift-image-registry
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: IBMCloudProviderSpec
    policies:
      - attributes:
          - name: serviceName
            value: cloud-object-storage
        roles:
          - crn:v1:bluemix:public:iam::::role:Viewer
          - crn:v1:bluemix:public:iam::::role:Operator
          - crn:v1:bluemix:public:iam::::role:Editor
          - crn:v1:bluemix:public:iam::::serviceRole:Reader
          - crn:v1:bluemix:public:iam::::serviceRole:Writer
      - attributes:
          - name: resourceType
            value: resource-group
        roles:
          - crn:v1:bluemix:public:iam::::role:Viewer

```

5. Create the service ID for each credential request, assign the policies defined, create an API key, and generate the secret:

```

$ ccoctl ibmcloud create-service-id \
  --credentials-requests-dir=<path_to_credential_requests_directory> \1
  --name=<cluster_name> \2
  --output-dir=<installation_directory> \3
  --resource-group-name=<resource_group_name> \4

```

- 1 Specify the directory containing the files for the component **CredentialsRequest** objects.
- 2 Specify the name of the OpenShift Container Platform cluster.
- 3 Optional: Specify the directory in which you want the **ccoctl** utility to create objects. By default, the utility creates objects in the directory in which the commands are run.
- 4 Optional: Specify the name of the resource group used for scoping the access policies.

**NOTE**

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

If an incorrect resource group name is provided, the installation fails during the bootstrap phase. To find the correct resource group name, run the following command:

```
$ grep resourceGroupName <installation_directory>/manifests/cluster-
infrastructure-02-config.yml
```

Verification

- Ensure that the appropriate secrets were generated in your cluster's **manifests** directory.

10.5.8. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.

**IMPORTANT**

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.

Procedure

- Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

- 1** For **<installation_directory>**, specify the location of your customized **./install-config.yaml** file.
- 2** To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to **<installation_directory>/openshift_install.log**.



IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

10.5.9. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.



IMPORTANT

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**.
To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.

**NOTE**

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your PATH.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

10.5.10. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

■

system:admin

Additional resources

- [Accessing the web console](#)

10.5.11. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- [About remote health monitoring](#)

10.5.12. Next steps

- [Customize your cluster](#).
- If necessary, you can [opt out of remote health reporting](#) .

10.6. INSTALLING A CLUSTER ON IBM CLOUD WITH NETWORK CUSTOMIZATIONS

In OpenShift Container Platform version 4.16, you can install a cluster with a customized network configuration on infrastructure that the installation program provisions on IBM Cloud®. By customizing your network configuration, your cluster can coexist with existing IP address allocations in your environment and integrate with existing MTU and VXLAN configurations. To customize the installation, you modify parameters in the **install-config.yaml** file before you install the cluster.

You must set most of the network configuration parameters during installation, and you can modify only **kubeProxy** configuration parameters in a running cluster.

10.6.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [configured an IBM Cloud® account](#) to host the cluster.
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.
- You configured the **ccoctl** utility before you installed the cluster. For more information, see [Configuring IAM for IBM Cloud®](#) .

10.6.2. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.



IMPORTANT

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

10.6.3. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.



NOTE

You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:


```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.



NOTE

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.



NOTE

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```



NOTE

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1 Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

10.6.4. Obtaining the installation program

Before you install OpenShift Container Platform, download the installation file on the host you are using for installation.

Prerequisites

- You have a computer that runs Linux or macOS, with 500 MB of local disk space.

Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select your infrastructure provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.



IMPORTANT

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.



IMPORTANT

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.

4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#) . This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

10.6.5. Exporting the API key

You must set the API key you created as a global variable; the installation program ingests the variable during startup to set the API key.

Prerequisites

- You have created either a user API key or service ID API key for your IBM Cloud® account.

Procedure

- Export your API key for your account as a global variable:

```
$ export IC_API_KEY=<api_key>
```



IMPORTANT

You must set the variable name exactly as specified; the installation program expects the variable name to be present during startup.

10.6.6. Creating the installation configuration file

You can customize the OpenShift Container Platform cluster you install on IBM Cloud®.

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create the **install-config.yaml** file.
 - a. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install create install-config --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the directory name to store the files that the installation program creates.

When specifying the directory:


- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
- Use an empty directory. Some installation assets, such as bootstrap X.509 certificates, have short expiration intervals, therefore you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

- b. At the prompts, provide the configuration details for your cloud:
 - i. Optional: Select an SSH key to use to access your cluster machines.

**NOTE**

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- ii. Select **ibmcloud** as the platform to target.
 - iii. Select the region to deploy the cluster to.
 - iv. Select the base domain to deploy the cluster to. The base domain corresponds to the public DNS zone that you created for your cluster.
 - v. Enter a descriptive name for your cluster.
2. Modify the **install-config.yaml** file. You can find more information about the available parameters in the "Installation configuration parameters" section.
 3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.

**IMPORTANT**

The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

Additional resources

- [Installation configuration parameters for IBM Cloud®](#)

10.6.6.1. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 10.4. Minimum resource requirements

Machine	Operating System	vCPU	Virtual RAM	Storage	Input/Output Per Second (IOPS)
Bootstrap	RHCOS	4	16 GB	100 GB	300
Control plane	RHCOS	4	16 GB	100 GB	300
Compute	RHCOS	2	8 GB	100 GB	300



NOTE

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

10.6.6.2. Tested instance types for IBM Cloud

The following IBM Cloud® instance types have been tested with OpenShift Container Platform.

Example 10.2. Machine series

- **c4.***
- **c5.***
- **c5a.***
- **i3.***
- **m4.***
- **m5.***
- **m5a.***
- **m6a.***
- **m6i.***
- **r4.***
- **r5.***
- **r5a.***
- **r6i.***
- **t3.***
- **t3a.***

Additional resources

- [Optimizing storage](#)

10.6.6.3. Sample customized install-config.yaml file for IBM Cloud

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.



IMPORTANT

This sample YAML file is provided for reference only. You must obtain your **install-config.yaml** file by using the installation program and then modify it.

```

apiVersion: v1
baseDomain: example.com 1
controlPlane: 2 3
  hyperthreading: Enabled 4
  name: master
  platform:
    ibmcloud: {}
  replicas: 3
compute: 5 6
- hyperthreading: Enabled 7
  name: worker
  platform:
    ibmcloud: {}
  replicas: 3
metadata:
  name: test-cluster 8
networking: 9
  clusterNetwork:
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  machineNetwork:
  - cidr: 10.0.0.0/16
  networkType: OVNKubernetes 10
  serviceNetwork:
  - 172.30.0.0/16
platform:
  ibmcloud:
    region: us-south 11
credentialsMode: Manual
publish: External
pullSecret: '{"auths": ...}' 12
fips: false 13
sshKey: ssh-ed25519 AAAA... 14

```

1 8 11 12 Required. The installation program prompts you for this value.

2 5 9 If you do not provide these parameters and values, the installation program provides the default value.

- 3 6 The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute**
- 4 7 Enables or disables simultaneous multithreading, also known as Hyper-Threading. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores. You can disable it by setting the parameter value to **Disabled**. If you disable simultaneous multithreading in some cluster machines, you must disable it in all cluster machines.



IMPORTANT

If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance. Use larger machine types, such as **n1-standard-8**, for your machines if you disable simultaneous multithreading.

- 10 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 13 Enables or disables FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.




IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 14 Optional: provide the **sshKey** value that you use to access the machines in your cluster.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

10.6.6.4. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```


- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when

http/https proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.



NOTE

The installation program does not support the proxy **readinessEndpoints** field.



NOTE

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.



NOTE

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

10.6.7. Manually creating IAM

Installing the cluster requires that the Cloud Credential Operator (CCO) operate in manual mode. While the installation program configures the CCO for manual mode, you must specify the identity and access management secrets for your cloud provider.

You can use the Cloud Credential Operator (CCO) utility (**ccoctl**) to create the required IBM Cloud® resources.

Prerequisites

- You have configured the **ccoctl** binary.
- You have an existing **install-config.yaml** file.

Procedure

1. Edit the **install-config.yaml** configuration file so that it contains the **credentialsMode** parameter set to **Manual**.

Example **install-config.yaml** configuration file

```
apiVersion: v1
baseDomain: cluster1.example.com
credentialsMode: Manual 1
```

```
compute:
- architecture: amd64
  hyperthreading: Enabled
```

- 1 This line is added to set the **credentialsMode** parameter to **Manual**.
2. To generate the manifests, run the following command from the directory that contains the installation program:

```
$ openshift-install create manifests --dir <installation_directory>
```

3. From the directory that contains the installation program, set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/' {print $3})
```

4. Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included \1 \
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \2 \
  --to=<path_to_directory_for_credentials_requests> \3
```

- 1 The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2 Specify the location of the **install-config.yaml** file.
- 3 Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

This command creates a YAML file for each **CredentialsRequest** object.

Sample **CredentialsRequest** object

```
apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  labels:
    controller-tools.k8s.io: "1.0"
  name: openshift-image-registry-ibmcos
  namespace: openshift-cloud-credential-operator
spec:
  secretRef:
    name: installer-cloud-credentials
    namespace: openshift-image-registry
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: IBMCloudProviderSpec
    policies:
```

```

- attributes:
  - name: serviceName
    value: cloud-object-storage
  roles:
  - crn:v1:bluemix:public:iam::::role:Viewer
  - crn:v1:bluemix:public:iam::::role:Operator
  - crn:v1:bluemix:public:iam::::role:Editor
  - crn:v1:bluemix:public:iam::::serviceRole:Reader
  - crn:v1:bluemix:public:iam::::serviceRole:Writer
- attributes:
  - name: resourceType
    value: resource-group
  roles:
  - crn:v1:bluemix:public:iam::::role:Viewer

```

5. Create the service ID for each credential request, assign the policies defined, create an API key, and generate the secret:

```

$ ccoctl ibmcloud create-service-id \
  --credentials-requests-dir=<path_to_credential_requests_directory> \ 1
  --name=<cluster_name> \ 2
  --output-dir=<installation_directory> \ 3
  --resource-group-name=<resource_group_name> \ 4

```

- 1** Specify the directory containing the files for the component **CredentialsRequest** objects.
- 2** Specify the name of the OpenShift Container Platform cluster.
- 3** Optional: Specify the directory in which you want the **ccoctl** utility to create objects. By default, the utility creates objects in the directory in which the commands are run.
- 4** Optional: Specify the name of the resource group used for scoping the access policies.

NOTE

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

If an incorrect resource group name is provided, the installation fails during the bootstrap phase. To find the correct resource group name, run the following command:

```

$ grep resourceGroupName <installation_directory>/manifests/cluster-
infrastructure-02-config.yml

```

Verification

- Ensure that the appropriate secrets were generated in your cluster's **manifests** directory.

10.6.8. Network configuration phases

There are two phases prior to OpenShift Container Platform installation where you can customize the network configuration.

Phase 1

You can customize the following network-related fields in the **install-config.yaml** file before you create the manifest files:

- **networking.networkType**
- **networking.clusterNetwork**
- **networking.serviceNetwork**
- **networking.machineNetwork**

For more information on these fields, refer to *Installation configuration parameters*.



NOTE

Set the **networking.machineNetwork** to match the CIDR that the preferred NIC resides in.



IMPORTANT

The CIDR range **172.17.0.0/16** is reserved by libVirt. You cannot use this range or any range that overlaps with this range for any networks in your cluster.


Phase 2

After creating the manifest files by running **openshift-install create manifests**, you can define a customized Cluster Network Operator manifest with only the fields you want to modify. You can use the manifest to specify advanced network configuration.

You cannot override the values specified in phase 1 in the **install-config.yaml** file during phase 2. However, you can further customize the network plugin during phase 2.

10.6.9. Specifying advanced network configuration

You can use advanced network configuration for your network plugin to integrate your cluster into your existing network environment. You can specify advanced network configuration only before you install the cluster.



IMPORTANT

Customizing your network configuration by modifying the OpenShift Container Platform manifest files created by the installation program is not supported. Applying a manifest file that you create, as in the following procedure, is supported.

Prerequisites

- You have created the **install-config.yaml** file and completed any modifications to it.

Procedure

1. Change to the directory that contains the installation program and create the manifests:

```
$ ./openshift-install create manifests --dir <installation_directory> 1
```

1 **<installation_directory>** specifies the name of the directory that contains the **install-config.yaml** file for your cluster.

2. Create a stub manifest file for the advanced network configuration that is named **cluster-network-03-config.yml** in the **<installation_directory>/manifests/** directory:

```
apiVersion: operator.openshift.io/v1
kind: Network
metadata:
  name: cluster
spec:
```

3. Specify the advanced network configuration for your cluster in the **cluster-network-03-config.yml** file, such as in the following example:

Enable IPsec for the OVN-Kubernetes network provider

```
apiVersion: operator.openshift.io/v1
kind: Network
metadata:
  name: cluster
spec:
  defaultNetwork:
    ovnKubernetesConfig:
      ipsecConfig:
        mode: Full
```

4. Optional: Back up the **manifests/cluster-network-03-config.yml** file. The installation program consumes the **manifests/** directory when you create the Ignition config files.

10.6.10. Cluster Network Operator configuration

The configuration for the cluster network is specified as part of the Cluster Network Operator (CNO) configuration and stored in a custom resource (CR) object that is named **cluster**. The CR specifies the fields for the **Network** API in the **operator.openshift.io** API group.

The CNO configuration inherits the following fields during cluster installation from the **Network** API in the **Network.config.openshift.io** API group:

clusterNetwork

IP address pools from which pod IP addresses are allocated.

serviceNetwork

IP address pool for services.

defaultNetwork.type

Cluster network plugin. **OVNKubernetes** is the only supported plugin during installation.

You can specify the cluster network plugin configuration for your cluster by setting the fields for the **defaultNetwork** object in the CNO object named **cluster**.

10.6.10.1. Cluster Network Operator configuration object

The fields for the Cluster Network Operator (CNO) are described in the following table:

Table 10.5. Cluster Network Operator configuration object


Field	Type	Description
metadata.name	string	The name of the CNO object. This name is always cluster .
spec.clusterNetwork	array	<p>A list specifying the blocks of IP addresses from which pod IP addresses are allocated and the subnet prefix length assigned to each individual node in the cluster. For example:</p> <pre>spec: clusterNetwork: - cidr: 10.128.0.0/19 hostPrefix: 23 - cidr: 10.128.32.0/19 hostPrefix: 23</pre>
spec.serviceNetwork	array	<p>A block of IP addresses for services. The OpenShift SDN and OVN-Kubernetes network plugins support only a single IP address block for the service network. For example:</p> <pre>spec: serviceNetwork: - 172.30.0.0/14</pre> <p>You can customize this field only in the install-config.yaml file before you create the manifests. The value is read-only in the manifest file.</p>
spec.defaultNetwork	object	Configures the network plugin for the cluster network.
spec.kubeProxy Config	object	The fields for this object specify the kube-proxy configuration. If you are using the OVN-Kubernetes cluster network plugin, the kube-proxy configuration has no effect.

defaultNetwork object configuration

The values for the **defaultNetwork** object are defined in the following table:

Table 10.6. defaultNetwork object

Field	Type	Description
-------	------	-------------

Field	Type	Description
type	string	<p>OVNKubernetes. The Red Hat OpenShift Networking network plugin is selected during installation. This value cannot be changed after cluster installation.</p>  <p>NOTE</p> <p>OpenShift Container Platform uses the OVN-Kubernetes network plugin by default. OpenShift SDN is no longer available as an installation choice for new clusters.</p>
ovnKubernetesConfig	object	This object is only valid for the OVN-Kubernetes network plugin.

Configuration for the OVN-Kubernetes network plugin

The following table describes the configuration fields for the OVN-Kubernetes network plugin:

Table 10.7. ovnKubernetesConfig object

Field	Type	Description
mtu	integer	<p>The maximum transmission unit (MTU) for the Geneve (Generic Network Virtualization Encapsulation) overlay network. This is detected automatically based on the MTU of the primary network interface. You do not normally need to override the detected MTU.</p> <p>If the auto-detected value is not what you expect it to be, confirm that the MTU on the primary network interface on your nodes is correct. You cannot use this option to change the MTU value of the primary network interface on the nodes.</p> <p>If your cluster requires different MTU values for different nodes, you must set this value to 100 less than the lowest MTU value in your cluster. For example, if some nodes in your cluster have an MTU of 9001, and some have an MTU of 1500, you must set this value to 1400.</p>
genevePort	integer	The port to use for all Geneve packets. The default value is 6081 . This value cannot be changed after cluster installation.
ipsecConfig	object	Specify a configuration object for customizing the IPsec configuration.
ipv4	object	Specifies a configuration object for IPv4 settings.
ipv6	object	Specifies a configuration object for IPv6 settings.


Field	Type	Description
policyAuditConfig	object	Specify a configuration object for customizing network policy audit logging. If unset, the defaults audit log settings are used.
gatewayConfig	object	Optional: Specify a configuration object for customizing how egress traffic is sent to the node gateway.  <p>NOTE</p> <p>While migrating egress traffic, you can expect some disruption to workloads and service traffic until the Cluster Network Operator (CNO) successfully rolls out the changes.</p>

Table 10.8. `ovnKubernetesConfig.ipv4` object

Field	Type	Description
internalTransitSwitchSubnet	string	If your existing network infrastructure overlaps with the 100.88.0.0/16 IPv4 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. The subnet for the distributed transit switch that enables east-west traffic. This subnet cannot overlap with any other subnets used by OVN-Kubernetes or on the host itself. It must be large enough to accommodate one IP address per node in your cluster. The default value is 100.88.0.0/16 .
internalJoinSubnet	string	If your existing network infrastructure overlaps with the 100.64.0.0/16 IPv4 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. You must ensure that the IP address range does not overlap with any other subnet used by your OpenShift Container Platform installation. The IP address range must be larger than the maximum number of nodes that can be added to the cluster. For example, if the clusterNetwork.cidr value is 10.128.0.0/14 and the clusterNetwork.hostPrefix value is /23 , then the maximum number of nodes is 2⁽²³⁻¹⁴⁾=512 . The default value is 100.64.0.0/16 .

Table 10.9. `ovnKubernetesConfig.ipv6` object

Field	Type	Description
-------	------	-------------

Field	Type	Description
internalTransitSwitchSubnet	string	<p>If your existing network infrastructure overlaps with the fd97::/64 IPv6 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. The subnet for the distributed transit switch that enables east-west traffic. This subnet cannot overlap with any other subnets used by OVN-Kubernetes or on the host itself. It must be large enough to accommodate one IP address per node in your cluster.</p> <p>The default value is fd97::/64.</p>
internalJoinSubnet	string	<p>If your existing network infrastructure overlaps with the fd98::/64 IPv6 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. You must ensure that the IP address range does not overlap with any other subnet used by your OpenShift Container Platform installation. The IP address range must be larger than the maximum number of nodes that can be added to the cluster.</p> <p>The default value is fd98::/64.</p>

Table 10.10. **policyAuditConfig** object

Field	Type	Description
rateLimit	integer	The maximum number of messages to generate every second per node. The default value is 20 messages per second.
maxFileSize	integer	The maximum size for the audit log in bytes. The default value is 50000000 or 50 MB.
maxLogFiles	integer	The maximum number of log files that are retained.
destination	string	<p>One of the following additional audit log targets:</p> <p>libc The libc syslog() function of the journald process on the host.</p> <p>udp:<host>:<port> A syslog server. Replace <host>:<port> with the host and port of the syslog server.</p> <p>unix:<file> A Unix Domain Socket file specified by <file>.</p> <p>null Do not send the audit logs to any additional target.</p>
syslogFacility	string	The syslog facility, such as kern , as defined by RFC5424. The default value is local0 .

Table 10.11. gatewayConfig object

Field	Type	Description
routingViaHost	boolean	<p>Set this field to true to send egress traffic from pods to the host networking stack. For highly-specialized installations and applications that rely on manually configured routes in the kernel routing table, you might want to route egress traffic to the host networking stack. By default, egress traffic is processed in OVN to exit the cluster and is not affected by specialized routes in the kernel routing table. The default value is false.</p> <p>This field has an interaction with the Open vSwitch hardware offloading feature. If you set this field to true, you do not receive the performance benefits of the offloading because egress traffic is processed by the host networking stack.</p>
ipForwarding	object	<p>You can control IP forwarding for all traffic on OVN-Kubernetes managed interfaces by using the ipForwarding specification in the Network resource. Specify Restricted to only allow IP forwarding for Kubernetes related traffic. Specify Global to allow forwarding of all IP traffic. For new installations, the default is Restricted. For updates to OpenShift Container Platform 4.14 or later, the default is Global.</p>
ipv4	object	<p>Optional: Specify an object to configure the internal OVN-Kubernetes masquerade address for host to service traffic for IPv4 addresses.</p>
ipv6	object	<p>Optional: Specify an object to configure the internal OVN-Kubernetes masquerade address for host to service traffic for IPv6 addresses.</p>

Table 10.12. gatewayConfig.ipv4 object

Field	Type	Description
internalMasqueradeSubnet	string	<p>The masquerade IPv4 addresses that are used internally to enable host to service traffic. The host is configured with these IP addresses as well as the shared gateway bridge interface. The default value is 169.254.169.0/29.</p>

Table 10.13. gatewayConfig.ipv6 object

Field	Type	Description
internalMasqueradeSubnet	string	<p>The masquerade IPv6 addresses that are used internally to enable host to service traffic. The host is configured with these IP addresses as well as the shared gateway bridge interface. The default value is fd69::/125.</p>

Table 10.14. `ipsecConfig` object


Field	Type	Description
<code>mode</code>	<code>string</code>	<p>Specifies the behavior of the IPsec implementation. Must be one of the following values:</p> <ul style="list-style-type: none"> ● Disabled: IPsec is not enabled on cluster nodes. ● External: IPsec is enabled for network traffic with external hosts. ● Full: IPsec is enabled for pod traffic and network traffic with external hosts.

Example OVN-Kubernetes configuration with IPsec enabled

```

defaultNetwork:
  type: OVNKubernetes
  ovnKubernetesConfig:
    mtu: 1400
    genevePort: 6081
    ipsecConfig:
      mode: Full

```




IMPORTANT

Using OVNKubernetes can lead to a stack exhaustion problem on IBM Power®.

`kubeProxyConfig` object configuration (OpenShiftSDN container network interface only)

The values for the `kubeProxyConfig` object are defined in the following table:

Table 10.15. `kubeProxyConfig` object

Field	Type	Description
<code>iptablesSyncPeriod</code>	<code>string</code>	<p>The refresh period for <code>iptables</code> rules. The default value is 30s. Valid suffixes include s, m, and h and are described in the Go time package documentation.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p>NOTE</p> <p>Because of performance improvements introduced in OpenShift Container Platform 4.3 and greater, adjusting the <code>iptablesSyncPeriod</code> parameter is no longer necessary.</p> </div> </div>

Field	Type	Description
proxyArguments.iptables-min-sync-period	array	<p>The minimum duration before refreshing iptables rules. This field ensures that the refresh does not happen too frequently. Valid suffixes include s, m, and h and are described in the Go time package. The default value is:</p> <pre>kubeProxyConfig: proxyArguments: iptables-min-sync-period: - 0s</pre>

10.6.11. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.

Procedure

- Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

- 1** For **<installation_directory>**, specify the location of your customized **./install-config.yaml** file.
- 2** To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to **<installation_directory>/openshift_install.log**.



IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

10.6.12. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.



IMPORTANT

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**.
To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:


```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.

**NOTE**

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your PATH.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

10.6.13. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
-
```

```
system:admin
```

Additional resources

- [Accessing the web console](#)

10.6.14. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- [About remote health monitoring](#)

10.6.15. Next steps

- [Customize your cluster](#).
- If necessary, you can [opt out of remote health reporting](#).

10.7. INSTALLING A CLUSTER ON IBM CLOUD INTO AN EXISTING VPC

In OpenShift Container Platform version 4.16, you can install a cluster into an existing Virtual Private Cloud (VPC) on IBM Cloud®. The installation program provisions the rest of the required infrastructure, which you can then further customize. To customize the installation, you modify parameters in the **install-config.yaml** file before you install the cluster.

10.7.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [configured an IBM Cloud® account](#) to host the cluster.
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.
- You configured the **ccoctl** utility before you installed the cluster. For more information, see [Configuring IAM for IBM Cloud®](#).

10.7.2. About using a custom VPC

In OpenShift Container Platform 4.16, you can deploy a cluster into the subnets of an existing IBM® Virtual Private Cloud (VPC). Deploying OpenShift Container Platform into an existing VPC can help you avoid limit constraints in new accounts or more easily abide by the operational constraints that your

company's guidelines set. If you cannot obtain the infrastructure creation permissions that are required to create the VPC yourself, use this installation option.

Because the installation program cannot know what other components are in your existing subnets, it cannot choose subnet CIDRs and so forth. You must configure networking for the subnets to which you will install the cluster.

10.7.2.1. Requirements for using your VPC

You must correctly configure the existing VPC and its subnets before you install the cluster. The installation program does not create the following components:

- NAT gateways
- Subnets
- Route tables
- VPC network

The installation program cannot:

- Subdivide network ranges for the cluster to use
- Set route tables for the subnets
- Set VPC options like DHCP



NOTE


The installation program requires that you use the cloud-provided DNS server. Using a custom DNS server is not supported and causes the installation to fail.

10.7.2.2. VPC validation

The VPC and all of the subnets must be in an existing resource group. The cluster is deployed to the existing VPC.

As part of the installation, specify the following in the **install-config.yaml** file:

- The name of the existing resource group that contains the VPC and subnets (**networkResourceGroupName**)
- The name of the existing VPC (**vpcName**)
- The subnets that were created for control plane machines and compute machines (**controlPlaneSubnets** and **computeSubnets**)



NOTE

Additional installer-provisioned cluster resources are deployed to a separate resource group (**resourceGroupName**). You can specify this resource group before installing the cluster. If undefined, a new resource group is created for the cluster.

To ensure that the subnets that you provide are suitable, the installation program confirms the following:

- All of the subnets that you specify exist.
- For each availability zone in the region, you specify:
 - One subnet for control plane machines.
 - One subnet for compute machines.
- The machine CIDR that you specified contains the subnets for the compute machines and control plane machines.

**NOTE**

Subnet IDs are not supported.

10.7.2.3. Isolation between clusters

If you deploy OpenShift Container Platform to an existing network, the isolation of cluster services is reduced in the following ways:

- You can install multiple OpenShift Container Platform clusters in the same VPC.
- ICMP ingress is allowed to the entire network.
- TCP port 22 ingress (SSH) is allowed to the entire network.
- Control plane TCP 6443 ingress (Kubernetes API) is allowed to the entire network.
- Control plane TCP 22623 ingress (MCS) is allowed to the entire network.

10.7.3. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.

**IMPORTANT**


If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

10.7.4. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.



NOTE

You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.



NOTE

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.

**NOTE**

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```

**NOTE**

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

10.7.5. Obtaining the installation program


Before you install OpenShift Container Platform, download the installation file on the host you are using for installation.

Prerequisites

- You have a computer that runs Linux or macOS, with 500 MB of local disk space.


Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select your infrastructure provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.



IMPORTANT

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.



IMPORTANT

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.

4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#) . This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

10.7.6. Exporting the API key

You must set the API key you created as a global variable; the installation program ingests the variable during startup to set the API key.


Prerequisites

- You have created either a user API key or service ID API key for your IBM Cloud® account.

Procedure

- Export your API key for your account as a global variable:

```
$ export IC_API_KEY=<api_key>
```



IMPORTANT

You must set the variable name exactly as specified; the installation program expects the variable name to be present during startup.

10.7.7. Creating the installation configuration file

You can customize the OpenShift Container Platform cluster you install on IBM Cloud®.

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create the **install-config.yaml** file.
 - a. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install create install-config --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the directory name to store the files that the installation program creates.

When specifying the directory:

- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
- Use an empty directory. Some installation assets, such as bootstrap X.509 certificates, have short expiration intervals, therefore you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

- b. At the prompts, provide the configuration details for your cloud:
 - i. Optional: Select an SSH key to use to access your cluster machines.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- ii. Select **ibmcloud** as the platform to target.
 - iii. Select the region to deploy the cluster to.
 - iv. Select the base domain to deploy the cluster to. The base domain corresponds to the public DNS zone that you created for your cluster.
 - v. Enter a descriptive name for your cluster.
2. Modify the **install-config.yaml** file. You can find more information about the available parameters in the "Installation configuration parameters" section.

3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

Additional resources

- [Installation configuration parameters for IBM Cloud®](#)

10.7.7.1. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 10.16. Minimum resource requirements

Machine	Operating System	vCPU	Virtual RAM	Storage	Input/Output Per Second (IOPS)
Bootstrap	RHCOS	4	16 GB	100 GB	300
Control plane	RHCOS	4	16 GB	100 GB	300
Compute	RHCOS	2	8 GB	100 GB	300



NOTE

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

10.7.7.2. Tested instance types for IBM Cloud

The following IBM Cloud® instance types have been tested with OpenShift Container Platform.

Example 10.3. Machine series

- **c4.***


- **c5.***
- **c5a.***
- **i3.***
- **m4.***
- **m5.***
- **m5a.***
- **m6a.***
- **m6i.***
- **r4.***
- **r5.***
- **r5a.***
- **r6i.***
- **t3.***
- **t3a.***

Additional resources

- [Optimizing storage](#)

10.7.7.3. Sample customized install-config.yaml file for IBM Cloud

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.



IMPORTANT

This sample YAML file is provided for reference only. You must obtain your **install-config.yaml** file by using the installation program and then modify it.

```

apiVersion: v1
baseDomain: example.com 1
controlPlane: 2 3
  hyperthreading: Enabled 4
  name: master
  platform:
    ibmcloud: {}
  replicas: 3
compute: 5 6
- hyperthreading: Enabled 7
  name: worker

```



```

platform:
  ibmcloud: {}
  replicas: 3
metadata:
  name: test-cluster 8
networking:
  clusterNetwork:
    - cidr: 10.128.0.0/14 9
    hostPrefix: 23
  machineNetwork:
    - cidr: 10.0.0.0/16
  networkType: OVNKubernetes 10
  serviceNetwork:
    - 172.30.0.0/16
platform:
  ibmcloud:
    region: eu-gb 11
  resourceGroupName: eu-gb-example-cluster-rg 12
  networkResourceGroupName: eu-gb-example-existing-network-rg 13
  vpcName: eu-gb-example-network-1 14
  controlPlaneSubnets: 15
    - eu-gb-example-network-1-cp-eu-gb-1
    - eu-gb-example-network-1-cp-eu-gb-2
    - eu-gb-example-network-1-cp-eu-gb-3
  computeSubnets: 16
    - eu-gb-example-network-1-compute-eu-gb-1
    - eu-gb-example-network-1-compute-eu-gb-2
    - eu-gb-example-network-1-compute-eu-gb-3
credentialsMode: Manual
publish: External
pullSecret: '{"auths": ...}' 17
fips: false 18
sshKey: ssh-ed25519 AAAA... 19

```

- 1 8 11 17 Required. The installation program prompts you for this value.
- 2 5 If you do not provide these parameters and values, the installation program provides the default value.
- 3 6 The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.
- 4 7 Enables or disables simultaneous multithreading, also known as Hyper-Threading. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores. You can disable it by setting the parameter value to **Disabled**. If you disable simultaneous multithreading in some cluster machines, you must disable it in all cluster machines.



IMPORTANT

If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance. Use larger machine types, such as **n1-standard-8**, for your machines if you disable simultaneous multithreading.

- 9 The machine CIDR must contain the subnets for the compute machines and control plane machines.
- 10 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 12 The name of an existing resource group. All installer-provisioned cluster resources are deployed to this resource group. If undefined, a new resource group is created for the cluster.
- 13 Specify the name of the resource group that contains the existing virtual private cloud (VPC). The existing VPC and subnets should be in this resource group. The cluster will be installed to this VPC.
- 14 Specify the name of an existing VPC.
- 15 Specify the name of the existing subnets to which to deploy the control plane machines. The subnets must belong to the VPC that you specified. Specify a subnet for each availability zone in the region.
- 16 Specify the name of the existing subnets to which to deploy the compute machines. The subnets must belong to the VPC that you specified. Specify a subnet for each availability zone in the region.
- 18 Enables or disables FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.



IMPORTANT

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 19 Optional: provide the **sshKey** value that you use to access the machines in your cluster.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

10.7.7.4. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the

trustedCA field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.

- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.



NOTE

The installation program does not support the proxy **readinessEndpoints** field.



NOTE

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.



NOTE

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

10.7.8. Manually creating IAM

Installing the cluster requires that the Cloud Credential Operator (CCO) operate in manual mode. While the installation program configures the CCO for manual mode, you must specify the identity and access management secrets for your cloud provider.

You can use the Cloud Credential Operator (CCO) utility (**ccoctl**) to create the required IBM Cloud® resources.

Prerequisites

- You have configured the **ccoctl** binary.
- You have an existing **install-config.yaml** file.

Procedure

1. Edit the **install-config.yaml** configuration file so that it contains the **credentialsMode** parameter set to **Manual**.

Example **install-config.yaml** configuration file

-

```

apiVersion: v1
baseDomain: cluster1.example.com
credentialsMode: Manual 1
compute:
- architecture: amd64
  hyperthreading: Enabled

```

1 This line is added to set the **credentialsMode** parameter to **Manual**.

- To generate the manifests, run the following command from the directory that contains the installation program:

```
$ openshift-install create manifests --dir <installation_directory>
```

- From the directory that contains the installation program, set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/' {print $3})
```

- Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```

$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included 1 \
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml 2 \
  --to=<path_to_directory_for_credentials_requests> 3

```

1 The **--included** parameter includes only the manifests that your specific cluster configuration requires.

2 Specify the location of the **install-config.yaml** file.

3 Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

This command creates a YAML file for each **CredentialsRequest** object.

Sample CredentialsRequest object

```

apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  labels:
    controller-tools.k8s.io: "1.0"
  name: openshift-image-registry-ibmcos
  namespace: openshift-cloud-credential-operator
spec:
  secretRef:
    name: installer-cloud-credentials
    namespace: openshift-image-registry

```

```

providerSpec:
  apiVersion: cloudcredential.openshift.io/v1
  kind: IBMCloudProviderSpec
  policies:
    - attributes:
        - name: serviceName
          value: cloud-object-storage
      roles:
        - crn:v1:bluemix:public:iam::::role:Viewer
        - crn:v1:bluemix:public:iam::::role:Operator
        - crn:v1:bluemix:public:iam::::role:Editor
        - crn:v1:bluemix:public:iam::::serviceRole:Reader
        - crn:v1:bluemix:public:iam::::serviceRole:Writer
    - attributes:
        - name: resourceType
          value: resource-group
      roles:
        - crn:v1:bluemix:public:iam::::role:Viewer

```

5. Create the service ID for each credential request, assign the policies defined, create an API key, and generate the secret:

```

$ ccoctl ibmcloud create-service-id \
  --credentials-requests-dir=<path_to_credential_requests_directory> \1
  --name=<cluster_name> \2
  --output-dir=<installation_directory> \3
  --resource-group-name=<resource_group_name> \4

```

- 1 Specify the directory containing the files for the component **CredentialsRequest** objects.
- 2 Specify the name of the OpenShift Container Platform cluster.
- 3 Optional: Specify the directory in which you want the **ccoctl** utility to create objects. By default, the utility creates objects in the directory in which the commands are run.
- 4 Optional: Specify the name of the resource group used for scoping the access policies.



NOTE

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

If an incorrect resource group name is provided, the installation fails during the bootstrap phase. To find the correct resource group name, run the following command:

```
$ grep resourceGroupName <installation_directory>/manifests/cluster-
infrastructure-02-config.yml
```

Verification

- Ensure that the appropriate secrets were generated in your cluster's **manifests** directory.

10.7.9. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.

Procedure

- Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

- 1** For **<installation_directory>**, specify the location of your customized **./install-config.yaml** file.
- 2** To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to **<installation_directory>/openshift_install.log**.



IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
```

```
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

10.7.10. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.



IMPORTANT

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**. To check your **PATH**, execute the following command:

```
$ echo $PATH
```


Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.



NOTE

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your PATH.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

10.7.11. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

Additional resources

- [Accessing the web console](#)

10.7.12. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- [About remote health monitoring](#)

10.7.13. Next steps

- [Customize your cluster.](#)
- Optional: [Opt out of remote health reporting](#).

10.8. INSTALLING A PRIVATE CLUSTER ON IBM CLOUD

In OpenShift Container Platform version 4.16, you can install a private cluster into an existing VPC. The installation program provisions the rest of the required infrastructure, which you can further customize. To customize the installation, you modify parameters in the **install-config.yaml** file before you install the cluster.


10.8.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [configured an IBM Cloud® account](#) to host the cluster.
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.
- You configured the **ccoctl** utility before you installed the cluster. For more information, see [Configuring IAM for IBM Cloud®](#).

10.8.2. Private clusters

You can deploy a private OpenShift Container Platform cluster that does not expose external endpoints. Private clusters are accessible from only an internal network and are not visible to the internet.

By default, OpenShift Container Platform is provisioned to use publicly-accessible DNS and endpoints. A private cluster sets the DNS, Ingress Controller, and API server to private when you deploy your cluster. This means that the cluster resources are only accessible from your internal network and are not visible to the internet.



IMPORTANT

If the cluster has any public subnets, load balancer services created by administrators might be publicly accessible. To ensure cluster security, verify that these services are explicitly annotated as private.

To deploy a private cluster, you must:

- Use existing networking that meets your requirements. Your cluster resources might be shared between other clusters on the network.

- Create a DNS zone using IBM Cloud® DNS Services and specify it as the base domain of the cluster. For more information, see "Using IBM Cloud® DNS Services to configure DNS resolution".
- Deploy from a machine that has access to:
 - The API services for the cloud to which you provision.
 - The hosts on the network that you provision.
 - The internet to obtain installation media.

You can use any machine that meets these access requirements and follows your company's guidelines. For example, this machine can be a bastion host on your cloud network or a machine that has access to the network through a VPN.

10.8.3. Private clusters in IBM Cloud

To create a private cluster on IBM Cloud®, you must provide an existing private VPC and subnets to host the cluster. The installation program must also be able to resolve the DNS records that the cluster requires. The installation program configures the Ingress Operator and API server for only internal traffic.

The cluster still requires access to internet to access the IBM Cloud® APIs.

The following items are not required or created when you install a private cluster:

- Public subnets
- Public network load balancers, which support public ingress
- A public DNS zone that matches the **baseDomain** for the cluster

The installation program does use the **baseDomain** that you specify to create a private DNS zone and the required records for the cluster. The cluster is configured so that the Operators do not create public records for the cluster and all cluster machines are placed in the private subnets that you specify.

10.8.3.1. Limitations

Private clusters on IBM Cloud® are subject only to the limitations associated with the existing VPC that was used for cluster deployment.

10.8.4. About using a custom VPC

In OpenShift Container Platform 4.16, you can deploy a cluster into the subnets of an existing IBM® Virtual Private Cloud (VPC). Deploying OpenShift Container Platform into an existing VPC can help you avoid limit constraints in new accounts or more easily abide by the operational constraints that your company's guidelines set. If you cannot obtain the infrastructure creation permissions that are required to create the VPC yourself, use this installation option.

Because the installation program cannot know what other components are in your existing subnets, it cannot choose subnet CIDRs and so forth. You must configure networking for the subnets to which you will install the cluster.

10.8.4.1. Requirements for using your VPC

You must correctly configure the existing VPC and its subnets before you install the cluster. The installation program does not create the following components:

- NAT gateways
- Subnets
- Route tables
- VPC network

The installation program cannot:

- Subdivide network ranges for the cluster to use
- Set route tables for the subnets
- Set VPC options like DHCP



NOTE

The installation program requires that you use the cloud-provided DNS server. Using a custom DNS server is not supported and causes the installation to fail.

10.8.4.2. VPC validation

The VPC and all of the subnets must be in an existing resource group. The cluster is deployed to the existing VPC.

As part of the installation, specify the following in the **install-config.yaml** file:

- The name of the existing resource group that contains the VPC and subnets (**networkResourceGroupName**)
- The name of the existing VPC (**vpcName**)
- The subnets that were created for control plane machines and compute machines (**controlPlaneSubnets** and **computeSubnets**)



NOTE

Additional installer-provisioned cluster resources are deployed to a separate resource group (**resourceGroupName**). You can specify this resource group before installing the cluster. If undefined, a new resource group is created for the cluster.

To ensure that the subnets that you provide are suitable, the installation program confirms the following:

- All of the subnets that you specify exist.
- For each availability zone in the region, you specify:
 - One subnet for control plane machines.
 - One subnet for compute machines.

- The machine CIDR that you specified contains the subnets for the compute machines and control plane machines.

**NOTE**

Subnet IDs are not supported.

10.8.4.3. Isolation between clusters

If you deploy OpenShift Container Platform to an existing network, the isolation of cluster services is reduced in the following ways:

- You can install multiple OpenShift Container Platform clusters in the same VPC.
- ICMP ingress is allowed to the entire network.
- TCP port 22 ingress (SSH) is allowed to the entire network.
- Control plane TCP 6443 ingress (Kubernetes API) is allowed to the entire network.
- Control plane TCP 22623 ingress (MCS) is allowed to the entire network.

10.8.5. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.

**IMPORTANT**

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

10.8.6. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The **./openshift-install gather** command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.



NOTE

You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.



NOTE

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.


2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the **./openshift-install gather** command.

**NOTE**

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```

**NOTE**

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

10.8.7. Obtaining the installation program

Before you install OpenShift Container Platform, download the installation file on a bastion host on your cloud network or a machine that has access to the to the network through a VPN.

For more information about private cluster installation requirements, see "Private clusters".


Prerequisites

- You have a machine that runs Linux, for example Red Hat Enterprise Linux 8, with 500 MB of local disk space.

Procedure


1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.

2. Select your infrastructure provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.



IMPORTANT

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.



IMPORTANT

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.

4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#) . This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

10.8.8. Exporting the API key

You must set the API key you created as a global variable; the installation program ingests the variable during startup to set the API key.


Prerequisites

- You have created either a user API key or service ID API key for your IBM Cloud® account.

Procedure

- Export your API key for your account as a global variable:

```
$ export IC_API_KEY=<api_key>
```



IMPORTANT

You must set the variable name exactly as specified; the installation program expects the variable name to be present during startup.

10.8.9. Manually creating the installation configuration file

Installing the cluster requires that you manually create the installation configuration file.

Prerequisites

- You have an SSH public key on your local machine to provide to the installation program. The key will be used for SSH authentication onto your cluster nodes for debugging and disaster recovery.
- You have obtained the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create an installation directory to store your required installation assets in:

```
$ mkdir <installation_directory>
```



IMPORTANT

You must create a directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.


2. Customize the sample **install-config.yaml** file template that is provided and save it in the **<installation_directory>**.



NOTE

You must name this configuration file **install-config.yaml**.

3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the next step of the installation process. You must back it up now.

Additional resources

- [Installation configuration parameters for IBM Cloud®](#)

10.8.9.1. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 10.17. Minimum resource requirements

Machine	Operating System	vCPU	Virtual RAM	Storage	Input/Output Per Second (IOPS)
Bootstrap	RHCOS	4	16 GB	100 GB	300
Control plane	RHCOS	4	16 GB	100 GB	300
Compute	RHCOS	2	8 GB	100 GB	300



NOTE

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

10.8.9.2. Tested instance types for IBM Cloud

The following IBM Cloud® instance types have been tested with OpenShift Container Platform.

Example 10.4. Machine series

- **c4.***
- **c5.***
- **c5a.***
- **i3.***
- **m4.***
- **m5.***
- **m5a.***

- **m6a.***
- **m6i.***
- **r4.***
- **r5.***
- **r5a.***
- **r6i.***
- **t3.***
- **t3a.***

10.8.9.3. Sample customized install-config.yaml file for IBM Cloud

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.



IMPORTANT

This sample YAML file is provided for reference only. You must obtain your **install-config.yaml** file by using the installation program and then modify it.

```

apiVersion: v1
baseDomain: example.com 1
controlPlane: 2 3
  hyperthreading: Enabled 4
  name: master
  platform:
    ibmcloud: {}
  replicas: 3
compute: 5 6
- hyperthreading: Enabled 7
  name: worker
  platform:
    ibmcloud: {}
  replicas: 3
metadata:
  name: test-cluster 8
networking:
  clusterNetwork:
  - cidr: 10.128.0.0/14 9
    hostPrefix: 23
  machineNetwork:
  - cidr: 10.0.0.0/16 10
    networkType: OVNKubernetes 11
  serviceNetwork:
  - 172.30.0.0/16
platform:
  ibmcloud:

```

```

region: eu-gb 12
resourceGroupName: eu-gb-example-cluster-rg 13
networkResourceGroupName: eu-gb-example-existing-network-rg 14
vpcName: eu-gb-example-network-1 15
controlPlaneSubnets: 16
  - eu-gb-example-network-1-cp-eu-gb-1
  - eu-gb-example-network-1-cp-eu-gb-2
  - eu-gb-example-network-1-cp-eu-gb-3
computeSubnets: 17
  - eu-gb-example-network-1-compute-eu-gb-1
  - eu-gb-example-network-1-compute-eu-gb-2
  - eu-gb-example-network-1-compute-eu-gb-3
credentialsMode: Manual
publish: Internal 18
pullSecret: '{"auths": ...}' 19
fips: false 20
sshKey: ssh-ed25519 AAAA... 21

```

1 8 12 19 Required.

2 5 If you do not provide these parameters and values, the installation program provides the default value.

3 6 The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.

4 7 Enables or disables simultaneous multithreading, also known as Hyper-Threading. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores. You can disable it by setting the parameter value to **Disabled**. If you disable simultaneous multithreading in some cluster machines, you must disable it in all cluster machines.



IMPORTANT

If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance. Use larger machine types, such as **n1-standard-8**, for your machines if you disable simultaneous multithreading.

9 The machine CIDR must contain the subnets for the compute machines and control plane machines.

10 The CIDR must contain the subnets defined in **platform.ibmcloud.controlPlaneSubnets** and **platform.ibmcloud.computeSubnets**.

11 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.

13 The name of an existing resource group. All installer-provisioned cluster resources are deployed to this resource group. If undefined, a new resource group is created for the cluster.

14 Specify the name of the resource group that contains the existing virtual private cloud (VPC). The existing VPC and subnets should be in this resource group. The cluster will be installed to this VPC.

- 15 Specify the name of an existing VPC.
- 16 Specify the name of the existing subnets to which to deploy the control plane machines. The subnets must belong to the VPC that you specified. Specify a subnet for each availability zone in the region.
- 17 Specify the name of the existing subnets to which to deploy the compute machines. The subnets must belong to the VPC that you specified. Specify a subnet for each availability zone in the region.
- 18 How to publish the user-facing endpoints of your cluster. Set **publish** to **Internal** to deploy a private cluster. The default value is **External**.
- 20 Enables or disables FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.



IMPORTANT

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 21 Optional: provide the **sshKey** value that you use to access the machines in your cluster.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

10.8.9.4. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.


For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.



NOTE

The installation program does not support the proxy **readinessEndpoints** field.


**NOTE**

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.

**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

10.8.10. Manually creating IAM

Installing the cluster requires that the Cloud Credential Operator (CCO) operate in manual mode. While the installation program configures the CCO for manual mode, you must specify the identity and access management secrets for your cloud provider.

You can use the Cloud Credential Operator (CCO) utility (**ccoctl**) to create the required IBM Cloud® resources.

Prerequisites

- You have configured the **ccoctl** binary.
- You have an existing **install-config.yaml** file.

Procedure

1. Edit the **install-config.yaml** configuration file so that it contains the **credentialsMode** parameter set to **Manual**.

Example **install-config.yaml** configuration file

```
apiVersion: v1
baseDomain: cluster1.example.com
credentialsMode: Manual 1
compute:
- architecture: amd64
  hyperthreading: Enabled
```

- 1** This line is added to set the **credentialsMode** parameter to **Manual**.

2. To generate the manifests, run the following command from the directory that contains the installation program:


```
$ openshift-install create manifests --dir <installation_directory>
```

- From the directory that contains the installation program, set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/' {print $3})
```

- Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included \ 1 \
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \ 2 \
  --to=<path_to_directory_for_credentials_requests> \ 3
```

- The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- Specify the location of the **install-config.yaml** file.
- Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

This command creates a YAML file for each **CredentialsRequest** object.

Sample CredentialsRequest object

```
apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  labels:
    controller-tools.k8s.io: "1.0"
  name: openshift-image-registry-ibmcos
  namespace: openshift-cloud-credential-operator
spec:
  secretRef:
    name: installer-cloud-credentials
    namespace: openshift-image-registry
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: IBMCloudProviderSpec
    policies:
      - attributes:
          - name: serviceName
            value: cloud-object-storage
      roles:
        - crn:v1:bluemix:public:iam::::role:Viewer
        - crn:v1:bluemix:public:iam::::role:Operator
        - crn:v1:bluemix:public:iam::::role:Editor
        - crn:v1:bluemix:public:iam::::serviceRole:Reader
        - crn:v1:bluemix:public:iam::::serviceRole:Writer
```

```

- attributes:
  - name: resourceType
    value: resource-group
  roles:
  - crn:v1:bluemix:public:iam:::role:Viewer

```


5. Create the service ID for each credential request, assign the policies defined, create an API key, and generate the secret:

```

$ ccoctl ibmcloud create-service-id \
  --credentials-requests-dir=<path_to_credential_requests_directory> \ 1
  --name=<cluster_name> \ 2
  --output-dir=<installation_directory> \ 3
  --resource-group-name=<resource_group_name> \ 4

```

- 1** Specify the directory containing the files for the component **CredentialsRequest** objects.
- 2** Specify the name of the OpenShift Container Platform cluster.
- 3** Optional: Specify the directory in which you want the **ccoctl** utility to create objects. By default, the utility creates objects in the directory in which the commands are run.
- 4** Optional: Specify the name of the resource group used for scoping the access policies.



NOTE

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

If an incorrect resource group name is provided, the installation fails during the bootstrap phase. To find the correct resource group name, run the following command:

```

$ grep resourceGroupName <installation_directory>/manifests/cluster-
infrastructure-02-config.yml

```

Verification

- Ensure that the appropriate secrets were generated in your cluster's **manifests** directory.

10.8.11. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.

Procedure

- Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

- 1** For **<installation_directory>**, specify the location of your customized **./install-config.yaml** file.
- 2** To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to **<installation_directory>/./openshift_install.log**.



IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

10.8.12. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.



IMPORTANT

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**. To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

-

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.



NOTE

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your **PATH**.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

10.8.13. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

Additional resources

- [Accessing the web console](#)

10.8.14. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- [About remote health monitoring](#)

10.8.15. Next steps

- [Customize your cluster.](#)
- If necessary, you can [opt out of remote health reporting](#) .

10.9. INSTALLING A CLUSTER ON IBM CLOUD IN A RESTRICTED NETWORK

In OpenShift Container Platform 4.16, you can install a cluster in a restricted network by creating an internal mirror of the installation release content that is accessible to an existing Virtual Private Cloud (VPC) on IBM Cloud®.

10.9.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You [configured an IBM Cloud account](#) to host the cluster.
- You have a container image registry that is accessible to the internet and your restricted network. The container image registry should mirror the contents of the OpenShift image registry and contain the installation media. For more information, see [Mirroring images for a disconnected installation using the oc-mirror plugin](#).
- You have an existing VPC on IBM Cloud® that meets the following requirements:
 - The VPC contains the mirror registry or has firewall rules or a peering connection to access the mirror registry that is hosted elsewhere.
 - The VPC can access IBM Cloud® service endpoints using a public endpoint. If network restrictions limit access to public service endpoints, evaluate those services for alternate endpoints that might be available. For more information see [Access to IBM service endpoints](#).

You cannot use the VPC that the installation program provisions by default.

- If you plan on configuring endpoint gateways to use IBM Cloud® Virtual Private Endpoints, consider the following requirements:
 - Endpoint gateway support is currently limited to the **us-east** and **us-south** regions.
 - The VPC must allow traffic to and from the endpoint gateways. You can use the VPC's default security group, or a new security group, to allow traffic on port 443. For more information, see [Allowing endpoint gateway traffic](#).
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.
- You configured the **ccoctl** utility before you installed the cluster. For more information, see [Configuring IAM for IBM Cloud VPC](#).

10.9.2. About installations in restricted networks

In OpenShift Container Platform 4.16, you can perform an installation that does not require an active connection to the internet to obtain software components. Restricted network installations can be completed using installer-provisioned infrastructure or user-provisioned infrastructure, depending on the cloud platform to which you are installing the cluster.

10.9.2.1. Required internet access and an installation host

You complete the installation using a bastion host or portable device that can access both the internet and your closed network. You must use a host with internet access to:

- Download the installation program, the OpenShift CLI (**oc**), and the CCO utility (**ccoctl**).
- Use the installation program to locate the Red Hat Enterprise Linux CoreOS (RHCOS) image and create the installation configuration file.
- Use **oc** to extract **ccoctl** from the CCO container image.
- Use **oc** and **ccoctl** to configure IAM for IBM Cloud®.

10.9.2.2. Access to a mirror registry

To complete a restricted network installation, you must create a registry that mirrors the contents of the OpenShift image registry and contains the installation media.

You can create this registry on a mirror host, which can access both the internet and your restricted network, or by using other methods that meet your organization's security restrictions.

For more information on mirroring images for a disconnected installation, see "Additional resources".

10.9.2.3. Access to IBM service endpoints

The installation program requires access to the following IBM Cloud® service endpoints:

- Cloud Object Storage
- DNS Services
- Global Search
- Global Tagging
- Identity Services
- Resource Controller
- Resource Manager
- VPC



NOTE

If you are specifying an IBM® Key Protect for IBM Cloud® root key as part of the installation process, the service endpoint for Key Protect is also required.

By default, the public endpoint is used to access the service. If network restrictions limit access to public service endpoints, you can override the default behavior.

Before deploying the cluster, you can update the installation configuration file (**install-config.yaml**) to specify the URI of an alternate service endpoint. For more information on usage, see "Additional resources".

10.9.2.4. Additional limits

Clusters in restricted networks have the following additional limitations and restrictions:

- The **ClusterVersion** status includes an **Unable to retrieve available updates** error.
- By default, you cannot use the contents of the Developer Catalog because you cannot access the required image stream tags.

Additional resources

- [Mirroring images for a disconnected installation using the oc-mirror plugin](#)
- [Additional IBM Cloud configuration parameters](#)

10.9.3. About using a custom VPC

In OpenShift Container Platform 4.16, you can deploy a cluster into the subnets of an existing IBM® Virtual Private Cloud (VPC). Deploying OpenShift Container Platform into an existing VPC can help you avoid limit constraints in new accounts or more easily abide by the operational constraints that your company's guidelines set. If you cannot obtain the infrastructure creation permissions that are required to create the VPC yourself, use this installation option.

Because the installation program cannot know what other components are in your existing subnets, it cannot choose subnet CIDRs and so forth. You must configure networking for the subnets to which you will install the cluster.

10.9.3.1. Requirements for using your VPC

You must correctly configure the existing VPC and its subnets before you install the cluster. The installation program does not create the following components:

- NAT gateways
- Subnets
- Route tables
- VPC network

The installation program cannot:

- Subdivide network ranges for the cluster to use
- Set route tables for the subnets
- Set VPC options like DHCP



NOTE

The installation program requires that you use the cloud-provided DNS server. Using a custom DNS server is not supported and causes the installation to fail.

10.9.3.2. VPC validation

The VPC and all of the subnets must be in an existing resource group. The cluster is deployed to the existing VPC.

As part of the installation, specify the following in the **install-config.yaml** file:

- The name of the existing resource group that contains the VPC and subnets (**networkResourceGroupName**)
- The name of the existing VPC (**vpcName**)
- The subnets that were created for control plane machines and compute machines (**controlPlaneSubnets** and **computeSubnets**)



NOTE

Additional installer-provisioned cluster resources are deployed to a separate resource group (**resourceGroupName**). You can specify this resource group before installing the cluster. If undefined, a new resource group is created for the cluster.

To ensure that the subnets that you provide are suitable, the installation program confirms the following:

- All of the subnets that you specify exist.
- For each availability zone in the region, you specify:
 - One subnet for control plane machines.
 - One subnet for compute machines.
- The machine CIDR that you specified contains the subnets for the compute machines and control plane machines.



NOTE

Subnet IDs are not supported.

10.9.3.3. Isolation between clusters

If you deploy OpenShift Container Platform to an existing network, the isolation of cluster services is reduced in the following ways:

- You can install multiple OpenShift Container Platform clusters in the same VPC.
- ICMP ingress is allowed to the entire network.
- TCP port 22 ingress (SSH) is allowed to the entire network.
- Control plane TCP 6443 ingress (Kubernetes API) is allowed to the entire network.
- Control plane TCP 22623 ingress (MCS) is allowed to the entire network.

10.9.3.4. Allowing endpoint gateway traffic

If you are using IBM Cloud® Virtual Private endpoints, your Virtual Private Cloud (VPC) must be configured to allow traffic to and from the endpoint gateways.

A VPC's default security group is configured to allow all outbound traffic to endpoint gateways. Therefore, the simplest way to allow traffic between your VPC and endpoint gateways is to modify the default security group to allow inbound traffic on port 443.



NOTE

If you choose to configure a new security group, the security group must be configured to allow both inbound and outbound traffic.

Prerequisites

- You have installed the IBM Cloud® Command Line Interface utility (**ibmcloud**).

Procedure

- Obtain the identifier for the default security group by running the following command:

```
$ DEFAULT_SG=$(ibmcloud is vpc <your_vpc_name> --output JSON | jq -r
'.default_security_group.id')
```

- Add a rule that allows inbound traffic on port 443 by running the following command:

```
$ ibmcloud is security-group-rule-add $DEFAULT_SG inbound tcp --remote 0.0.0.0/0 --port-
min 443 --port-max 443
```



NOTE


Be sure that your endpoint gateways are configured to use this security group.

10.9.4. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The **./openshift-install gather** command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.

**NOTE**

You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.

**NOTE**

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.

**NOTE**

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```

**NOTE**

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

10.9.5. Exporting the API key

You must set the API key you created as a global variable; the installation program ingests the variable during startup to set the API key.

Prerequisites

- You have created either a user API key or service ID API key for your IBM Cloud® account.

Procedure

- Export your API key for your account as a global variable:

```
$ export IC_API_KEY=<api_key>
```

**IMPORTANT**

You must set the variable name exactly as specified; the installation program expects the variable name to be present during startup.

10.9.6. Downloading the RHCOS cluster image

The installation program requires the Red Hat Enterprise Linux CoreOS (RHCOS) image to install the cluster. While optional, downloading the Red Hat Enterprise Linux CoreOS (RHCOS) before deploying removes the need for internet access when creating the cluster.

Use the installation program to locate and download the Red Hat Enterprise Linux CoreOS (RHCOS) image.

Prerequisites

- The host running the installation program has internet access.

Procedure

1. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install coreos print-stream-json
```

2. Use the output of the command to find the location of the IBM Cloud® image.

```
.Example output
```

```
----  
"release": "415.92.202311241643-0",  
"formats": {  
  "qcow2.gz": {  
    "disk": {  
      "location": "https://rhcos.mirror.openshift.com/art/storage/prod/streams/4.15-  
9.2/builds/415.92.202311241643-0/x86_64/rhcos-415.92.202311241643-0-  
ibmcloud.x86_64.qcow2.gz",  
      "sha256":  
"6b562dee8431bec3b93adeac1cfefcd5e812d41e3b7d78d3e28319870ffc9eae",  
      "uncompressed-sha256":  
"5a0f9479505e525a30367b6a6a6547c86a8f03136f453c1da035f3aa5daa8bc9"  
----
```

3. Download and extract the image archive. Make the image available on the host that the installation program uses to create the cluster.

10.9.7. Manually creating the installation configuration file

Installing the cluster requires that you manually create the installation configuration file.

Prerequisites

- You have obtained the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have the **imageContentSourcePolicy.yaml** file that was created when you mirrored your registry.
- You have obtained the contents of the certificate for your mirror registry.

Procedure

1. Create an installation directory to store your required installation assets in:

```
$ mkdir <installation_directory>
```



IMPORTANT

You must create a directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

2. Customize the sample **install-config.yaml** file template that is provided and save it in the **<installation_directory>**.



NOTE

You must name this configuration file **install-config.yaml**.

When customizing the sample template, be sure to provide the information that is required for an installation in a restricted network:

- a. Update the **pullSecret** value to contain the authentication information for your registry:

```
pullSecret: '{"auths":{"<mirror_host_name>:5000": {"auth": "<credentials>","email":
"you@example.com"}}}'
```

For **<mirror_host_name>**, specify the registry domain name that you specified in the certificate for your mirror registry, and for **<credentials>**, specify the base64-encoded user name and password for your mirror registry.

- b. Add the **additionalTrustBundle** parameter and value.

```
additionalTrustBundle: |
  -----BEGIN CERTIFICATE-----
  /-----END CERTIFICATE-----
```

The value must be the contents of the certificate file that you used for your mirror registry. The certificate file can be an existing, trusted certificate authority, or the self-signed certificate that you generated for the mirror registry.

- c. Define the network and subnets for the VPC to install the cluster in under the parent **platform.ibmcloud** field:

```
vpcName: <existing_vpc>
controlPlaneSubnets: <control_plane_subnet>
computeSubnets: <compute_subnet>
```


For **platform.ibmcloud.vpcName**, specify the name for the existing IBM Cloud VPC. For **platform.ibmcloud.controlPlaneSubnets** and **platform.ibmcloud.computeSubnets**, specify the existing subnets to deploy the control plane machines and compute machines, respectively.

- d. Add the image content resources, which resemble the following YAML excerpt:

```
imageContentSources:
- mirrors:
  - <mirror_host_name>:5000/<repo_name>/release
    source: quay.io/openshift-release-dev/ocp-release
- mirrors:
  - <mirror_host_name>:5000/<repo_name>/release
    source: registry.redhat.io/ocp/release
```

For these values, use the **imageContentSourcePolicy.yaml** file that was created when you mirrored the registry.

- e. If network restrictions limit the use of public endpoints to access the required IBM Cloud® services, add the **serviceEndpoints** stanza to **platform.ibmcloud** to specify an alternate service endpoint.



NOTE

You can specify only one alternate service endpoint for each service.

Example of using alternate services endpoints

```
# ...
serviceEndpoints:
- name: IAM
  url: <iam_alternate_endpoint_url>
- name: VPC
  url: <vpc_alternate_endpoint_url>
- name: ResourceController
  url: <resource_controller_alternate_endpoint_url>
- name: ResourceManager
  url: <resource_manager_alternate_endpoint_url>
- name: DNSServices
  url: <dns_services_alternate_endpoint_url>
- name: COS
  url: <cos_alternate_endpoint_url>
- name: GlobalSearch
  url: <global_search_alternate_endpoint_url>
- name: GlobalTagging
  url: <global_tagging_alternate_endpoint_url>
# ...
```

- f. Optional: Set the publishing strategy to **Internal**:

```
publish: Internal
```

By setting this option, you create an internal Ingress Controller and a private load balancer.



NOTE

If you use the default value of **External**, your network must be able to access the public endpoint for IBM Cloud® Internet Services (CIS). CIS is not enabled for Virtual Private Endpoints.

3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the next step of the installation process. You must back it up now.

10.9.7.1. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.

- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with `.` to match subdomains only. For
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.

**NOTE**

The installation program does not support the proxy **readinessEndpoints** field.

**NOTE**

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.

**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

Additional resources


- [Installation configuration parameters for IBM Cloud®](#)

10.9.7.2. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 10.18. Minimum resource requirements

Machine	Operating System	vCPU	Virtual RAM	Storage	Input/Output Per Second (IOPS)
Bootstrap	RHCOS	4	16 GB	100 GB	300
Control plane	RHCOS	4	16 GB	100 GB	300
Compute	RHCOS	2	8 GB	100 GB	300



NOTE

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

10.9.7.3. Tested instance types for IBM Cloud

The following IBM Cloud® instance types have been tested with OpenShift Container Platform.

Example 10.5. Machine series

- **c4.***
- **c5.***
- **c5a.***
- **i3.***
- **m4.***
- **m5.***
- **m5a.***
- **m6a.***
- **m6i.***

- r4.*
- r5.*
- r5a.*
- r6i.*
- t3.*
- t3a.*

10.9.7.4. Sample customized install-config.yaml file for IBM Cloud

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.



IMPORTANT

This sample YAML file is provided for reference only. You must obtain your **install-config.yaml** file by using the installation program and then modify it.

```

apiVersion: v1
baseDomain: example.com 1
controlPlane: 2 3
  hyperthreading: Enabled 4
  name: master
  platform:
    ibm-cloud: {}
  replicas: 3
compute: 5 6
- hyperthreading: Enabled 7
  name: worker
  platform:
    ibmcloud: {}
  replicas: 3
metadata:
  name: test-cluster 8
networking:
  clusterNetwork:
  - cidr: 10.128.0.0/14 9
    hostPrefix: 23
  machineNetwork:
  - cidr: 10.0.0.0/16 10
  networkType: OVNKubernetes 11
  serviceNetwork:
  - 172.30.0.0/16
platform:
  ibmcloud:
    region: us-east 12
    resourceName: us-east-example-cluster-rg 13
    serviceEndpoints: 14

```

```

- name: IAM
  url: https://private.us-east.iam.cloud.ibm.com
- name: VPC
  url: https://us-east.private.iaas.cloud.ibm.com/v1
- name: ResourceController
  url: https://private.us-east.resource-controller.cloud.ibm.com
- name: ResourceManager
  url: https://private.us-east.resource-controller.cloud.ibm.com
- name: DNSServices
  url: https://api.private.dns-svcs.cloud.ibm.com/v1
- name: COS
  url: https://s3.direct.us-east.cloud-object-storage.appdomain.cloud
- name: GlobalSearch
  url: https://api.private.global-search-tagging.cloud.ibm.com
- name: GlobalTagging
  url: https://tags.private.global-search-tagging.cloud.ibm.com
networkResourceGroupName: us-east-example-existing-network-rg 15
vpcName: us-east-example-network-1 16
controlPlaneSubnets: 17
  - us-east-example-network-1-cp-us-east-1
  - us-east-example-network-1-cp-us-east-2
  - us-east-example-network-1-cp-us-east-3
computeSubnets: 18
  - us-east-example-network-1-compute-us-east-1
  - us-east-example-network-1-compute-us-east-2
  - us-east-example-network-1-compute-us-east-3
credentialsMode: Manual
pullSecret: '{"auths":{"<local_registry>":{"auth": "<credentials>","email": "you@example.com"}}}' 19
fips: false 20
sshKey: ssh-ed25519 AAAA... 21
additionalTrustBundle: | 22
  -----BEGIN CERTIFICATE-----
  <MY_TRUSTED_CA_CERT>
  -----END CERTIFICATE-----
imageContentSources: 23
- mirrors:
  - <local_registry>/<local_repository_name>/release
  source: quay.io/openshift-release-dev/ocp-release
- mirrors:
  - <local_registry>/<local_repository_name>/release
  source: quay.io/openshift-release-dev/ocp-v4.0-art-dev

```

1 8 12 Required.

2 5 If you do not provide these parameters and values, the installation program provides the default value.

3 6 The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.

4 7 Enables or disables simultaneous multithreading, also known as Hyper-Threading. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores. You can disable it by setting the parameter value to **Disabled**. If you disable simultaneous

multithreading in some cluster machines, you must disable it in all cluster machines.



IMPORTANT

If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance. Use larger machine types, such as **n1-standard-8**, for your machines if you disable simultaneous multithreading.

- 9 The machine CIDR must contain the subnets for the compute machines and control plane machines.
- 10 The CIDR must contain the subnets defined in **platform.ibmcloud.controlPlaneSubnets** and **platform.ibmcloud.computeSubnets**.
- 11 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 13 The name of an existing resource group. All installer-provisioned cluster resources are deployed to this resource group. If undefined, a new resource group is created for the cluster.
- 14 Based on the network restrictions of the VPC, specify alternate service endpoints as needed. This overrides the default public endpoint for the service.
- 15 Specify the name of the resource group that contains the existing virtual private cloud (VPC). The existing VPC and subnets should be in this resource group. The cluster will be installed to this VPC.
- 16 Specify the name of an existing VPC.
- 17 Specify the name of the existing subnets to which to deploy the control plane machines. The subnets must belong to the VPC that you specified. Specify a subnet for each availability zone in the region.
- 18 Specify the name of the existing subnets to which to deploy the compute machines. The subnets must belong to the VPC that you specified. Specify a subnet for each availability zone in the region.
- 19 For **<local_registry>**, specify the registry domain name, and optionally the port, that your mirror registry uses to serve content. For example, `registry.example.com` or `registry.example.com:5000`. For **<credentials>**, specify the base64-encoded user name and password for your mirror registry.
- 20 Enables or disables FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.



IMPORTANT

The use of FIPS Validated or Modules in Process cryptographic libraries is only supported on OpenShift Container Platform deployments on the **x86_64** architecture.

- 21 Optional: provide the **sshKey** value that you use to access the machines in your cluster.
- 22 Provide the contents of the certificate file that you used for your mirror registry.

- 23** Provide these values from the **metadata.name: release-0** section of the **imageContentSourcePolicy.yaml** file that was created when you mirrored the registry.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

10.9.8. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.



IMPORTANT

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**. To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:


```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.



NOTE

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your PATH.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

10.9.9. Manually creating IAM

Installing the cluster requires that the Cloud Credential Operator (CCO) operate in manual mode. While the installation program configures the CCO for manual mode, you must specify the identity and access management secrets for your cloud provider.

You can use the Cloud Credential Operator (CCO) utility (**ccoctl**) to create the required IBM Cloud® resources.

Prerequisites

- You have configured the **ccoctl** binary.
- You have an existing **install-config.yaml** file.

Procedure

1. Edit the **install-config.yaml** configuration file so that it contains the **credentialsMode** parameter set to **Manual**.

Example **install-config.yaml** configuration file

```
apiVersion: v1
baseDomain: cluster1.example.com
credentialsMode: Manual 1
compute:
- architecture: amd64
  hyperthreading: Enabled
```

- 1** This line is added to set the **credentialsMode** parameter to **Manual**.

2. To generate the manifests, run the following command from the directory that contains the installation program:

```
$ openshift-install create manifests --dir <installation_directory>
```

3. From the directory that contains the installation program, set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

4. Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included 1 \
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml 2 \
  --to=<path_to_directory_for_credentials_requests> 3
```

- 1** The **--included** parameter includes only the manifests that your specific cluster configuration requires.

- 2 Specify the location of the **install-config.yaml** file.
- 3 Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

This command creates a YAML file for each **CredentialsRequest** object.

Sample **CredentialsRequest** object

```

apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  labels:
    controller-tools.k8s.io: "1.0"
  name: openshift-image-registry-ibmcos
  namespace: openshift-cloud-credential-operator
spec:
  secretRef:
    name: installer-cloud-credentials
    namespace: openshift-image-registry
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: IBMCloudProviderSpec
    policies:
      - attributes:
          - name: serviceName
            value: cloud-object-storage
        roles:
          - crn:v1:bluemix:public:iam:::role:Viewer
          - crn:v1:bluemix:public:iam:::role:Operator
          - crn:v1:bluemix:public:iam:::role:Editor
          - crn:v1:bluemix:public:iam:::serviceRole:Reader
          - crn:v1:bluemix:public:iam:::serviceRole:Writer
      - attributes:
          - name: resourceType
            value: resource-group
        roles:
          - crn:v1:bluemix:public:iam:::role:Viewer

```

5. Create the service ID for each credential request, assign the policies defined, create an API key, and generate the secret:

```

$ ccoctl ibmcloud create-service-id \
  --credentials-requests-dir=<path_to_credential_requests_directory> \1
  --name=<cluster_name> \2
  --output-dir=<installation_directory> \3
  --resource-group-name=<resource_group_name> \4

```

- 1 Specify the directory containing the files for the component **CredentialsRequest** objects.
- 2 Specify the name of the OpenShift Container Platform cluster.
- 3 Optional: Specify the directory in which you want the **ccoctl** utility to create objects. By default, the utility creates objects in the directory in which the commands are run.

- 4 Optional: Specify the name of the resource group used for scoping the access policies.



NOTE

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

If an incorrect resource group name is provided, the installation fails during the bootstrap phase. To find the correct resource group name, run the following command:


```
$ grep resourceGroupName <installation_directory>/manifests/cluster-
infrastructure-02-config.yml
```

Verification

- Ensure that the appropriate secrets were generated in your cluster's **manifests** directory.

10.9.10. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
If the Red Hat Enterprise Linux CoreOS (RHCOS) image is available locally, the host running the installation program does not require internet access.
- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.

Procedure

1. Export the **OPENSIFT_INSTALL_OS_IMAGE_OVERRIDE** variable to specify the location of the Red Hat Enterprise Linux CoreOS (RHCOS) image by running the following command:

```
$ export OPENSIFT_INSTALL_OS_IMAGE_OVERRIDE="<path_to_image>/rhcos-
<image_version>-ibmcloud.x86_64.qcow2.gz"
```

2. Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

- 1** For **<installation_directory>**, specify the location of your customized **./install-config.yaml** file.
- 2** To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to **<installation_directory>/openshift_install.log**.



IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

10.9.11. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the

correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

Additional resources

- [Accessing the web console](#)

10.9.12. Post installation

Complete the following steps to complete the configuration of your cluster.

10.9.12.1. Disabling the default OperatorHub catalog sources

Operator catalogs that source content provided by Red Hat and community projects are configured for OperatorHub by default during an OpenShift Container Platform installation. In a restricted network environment, you must disable the default catalogs as a cluster administrator.

Procedure

- Disable the sources for the default catalogs by adding **disableAllDefaultSources: true** to the **OperatorHub** object:

```
$ oc patch OperatorHub cluster --type json \  
-p '[{"op": "add", "path": "/spec/disableAllDefaultSources", "value": true}]'
```

TIP

Alternatively, you can use the web console to manage catalog sources. From the **Administration** → **Cluster Settings** → **Configuration** → **OperatorHub** page, click the **Sources** tab, where you can create, update, delete, disable, and enable individual sources.

10.9.12.2. Installing the policy resources into the cluster

Mirroring the OpenShift Container Platform content using the oc-mirror OpenShift CLI (oc) plugin creates resources, which include **catalogSource-certified-operator-index.yaml** and **imageContentSourcePolicy.yaml**.

- The **ImageContentSourcePolicy** resource associates the mirror registry with the source registry and redirects image pull requests from the online registries to the mirror registry.
- The **CatalogSource** resource is used by Operator Lifecycle Manager (OLM) to retrieve information about the available Operators in the mirror registry, which lets users discover and install Operators.

After you install the cluster, you must install these resources into the cluster.

Prerequisites

- You have mirrored the image set to the registry mirror in the disconnected environment.
- You have access to the cluster as a user with the **cluster-admin** role.

Procedure

1. Log in to the OpenShift CLI as a user with the **cluster-admin** role.
2. Apply the YAML files from the results directory to the cluster:

```
$ oc apply -f ./oc-mirror-workspace/results-<id>/
```

Verification

1. Verify that the **ImageContentSourcePolicy** resources were successfully installed:

```
$ oc get imagecontentsourcepolicy
```

2. Verify that the **CatalogSource** resources were successfully installed:

```
$ oc get catalogsource --all-namespaces
```

10.9.13. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- [About remote health monitoring](#)

10.9.14. Next steps


- [Customize your cluster.](#)
- Optional: [Opt out of remote health reporting.](#)

10.10. INSTALLATION CONFIGURATION PARAMETERS FOR IBM CLOUD

Before you deploy an OpenShift Container Platform cluster on IBM Cloud®, you provide parameters to customize your cluster and the platform that hosts it. When you create the **install-config.yaml** file, you provide values for the required parameters through the command line. You can then modify the **install-config.yaml** file to customize your cluster further.

10.10.1. Available installation configuration parameters for IBM Cloud

The following tables specify the required, optional, and IBM Cloud-specific installation configuration parameters that you can set as part of the installation process.



NOTE

After installation, you cannot modify these parameters in the **install-config.yaml** file.

10.10.1.1. Required configuration parameters

Required installation configuration parameters are described in the following table:

Table 10.19. Required parameters

Parameter	Description	Values
apiVersion:	The API version for the install-config.yaml content. The current version is v1 . The installation program may also support older API versions.	String

Parameter	Description	Values
baseDomain:	The base domain of your cloud provider. The base domain is used to create routes to your OpenShift Container Platform cluster components. The full DNS name for your cluster is a combination of the baseDomain and metadata.name parameter values that uses the <metadata.name>.<baseDomain> format.	A fully-qualified domain or subdomain name, such as example.com .
metadata:	Kubernetes resource ObjectMeta , from which only the name parameter is consumed.	Object
metadata: name:	The name of the cluster. DNS records for the cluster are all subdomains of {{.metadata.name}} . {{.baseDomain}} .	String of lowercase letters, hyphens (-), and periods (.), such as dev .
platform:	The configuration for the specific platform upon which to perform the installation: aws, baremetal, azure, gcp, ibmcloud, nutanix, openstack, powervs, vsphere , or {} . For additional information about platform . <platform> parameters, consult the table for your specific platform that follows.	Object

Parameter	Description	Values
<code>pullSecret:</code>	Get a pull secret from Red Hat OpenShift Cluster Manager to authenticate downloading container images for OpenShift Container Platform components from services such as Quay.io.	<pre>{ "auths":{ "cloud.openshift.com":{ "auth":"b3Blb=", "email":"you@example.com" }, "quay.io":{ "auth":"b3Blb=", "email":"you@example.com" } } }</pre>

10.10.1.2. Network configuration parameters

You can customize your installation configuration based on the requirements of your existing network infrastructure. For example, you can expand the IP address block for the cluster network or provide different IP address blocks than the defaults.


Only IPv4 addresses are supported.




NOTE

Globalnet is not supported with Red Hat OpenShift Data Foundation disaster recovery solutions. For regional disaster recovery scenarios, ensure that you use a nonoverlapping range of private IP addresses for the cluster and service networks in each cluster.

Table 10.20. Network parameters

Parameter	Description	Values
<code>networking:</code>	The configuration for the cluster network.	Object  NOTE You cannot modify parameters specified by the networking object after installation.
<code>networking: networkType:</code>	The Red Hat OpenShift Networking network plugin to install.	OVNKubernetes. OVNKubernetes is a CNI plugin for Linux networks and hybrid networks that contain both Linux and Windows servers. The default value is OVNKubernetes .

Parameter	Description	Values
<code>networking: clusterNetwork:</code>	<p>The IP address blocks for pods.</p> <p>The default value is 10.128.0.0/14 with a host prefix of /23.</p> <p>If you specify multiple IP address blocks, the blocks must not overlap.</p>	<p>An array of objects. For example:</p> <pre>networking: clusterNetwork: - cidr: 10.128.0.0/14 hostPrefix: 23</pre>
<code>networking: clusterNetwork: cidr:</code>	<p>Required if you use networking.clusterNetwork. An IP address block.</p> <p>An IPv4 network.</p>	<p>An IP address block in Classless Inter-Domain Routing (CIDR) notation. The prefix length for an IPv4 block is between 0 and 32.</p>
<code>networking: clusterNetwork: hostPrefix:</code>	<p>The subnet prefix length to assign to each individual node. For example, if hostPrefix is set to 23 then each node is assigned a /23 subnet out of the given cidr. A hostPrefix value of 23 provides 510 ($2^{(32 - 23)} - 2$) pod IP addresses.</p>	<p>A subnet prefix.</p> <p>The default value is 23.</p>
<code>networking: serviceNetwork:</code>	<p>The IP address block for services. The default value is 172.30.0.0/16.</p> <p>The OVN-Kubernetes network plugins supports only a single IP address block for the service network.</p>	<p>An array with an IP address block in CIDR format. For example:</p> <pre>networking: serviceNetwork: - 172.30.0.0/16</pre>
<code>networking: machineNetwork:</code>	<p>The IP address blocks for machines.</p> <p>If you specify multiple IP address blocks, the blocks must not overlap.</p>	<p>An array of objects. For example:</p> <pre>networking: machineNetwork: - cidr: 10.0.0.0/16</pre>


Parameter	Description	Values
networking: machineNetwork: cidr:	Required if you use networking.machineNetwork . An IP address block. The default value is 10.0.0.0/16 for all platforms other than libvirt and IBM Power® Virtual Server. For libvirt, the default value is 192.168.126.0/24 . For IBM Power® Virtual Server, the default value is 192.168.0.0/24 . If you are deploying the cluster to an existing Virtual Private Cloud (VPC), the CIDR must contain the subnets defined in platform.ibmcloud.controlPlaneSubnets and platform.ibmcloud.computeSubnets .	An IP network block in CIDR notation. For example, 10.0.0.0/16 .  <div style="margin-left: 20px;"> <p>NOTE</p> <p>Set the networking.machineNetwork to match the CIDR that the preferred NIC resides in.</p> </div>

10.10.1.3. Optional configuration parameters


Optional installation configuration parameters are described in the following table:

Table 10.21. Optional parameters

Parameter	Description	Values
additionalTrustBundle:	A PEM-encoded X.509 certificate bundle that is added to the nodes' trusted certificate store. This trust bundle may also be used when a proxy has been configured.	String
capabilities:	Controls the installation of optional core cluster components. You can reduce the footprint of your OpenShift Container Platform cluster by disabling optional components. For more information, see the "Cluster capabilities" page in <i>Installing</i> .	String array
capabilities: baselineCapabilitySet:	Selects an initial set of optional capabilities to enable. Valid values are None , v4.11 , v4.12 and vCurrent . The default value is vCurrent .	String

Parameter	Description	Values
capabilities: additionalEnabledCapabilities:	Extends the set of optional capabilities beyond what you specify in baselineCapabilitySet . You may specify multiple capabilities in this parameter.	String array
cpuPartitioningMode:	Enables workload partitioning, which isolates OpenShift Container Platform services, cluster management workloads, and infrastructure pods to run on a reserved set of CPUs. Workload partitioning can only be enabled during installation and cannot be disabled after installation. While this field enables workload partitioning, it does not configure workloads to use specific CPUs. For more information, see the <i>Workload partitioning</i> page in the <i>Scalability and Performance</i> section.	None or AllNodes . None is the default value.
compute:	The configuration for the machines that comprise the compute nodes.	Array of MachinePool objects.
compute: architecture:	Determines the instruction set architecture of the machines in the pool. Currently, clusters with varied architectures are not supported. All pools must specify the same architecture. Valid values are amd64 (the default).	String
compute: hyperthreading:	Whether to enable or disable simultaneous multithreading, or hyperthreading , on compute machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.  IMPORTANT If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.	Enabled or Disabled

Parameter	Description	Values
<code>compute: name:</code>	Required if you use compute . The name of the machine pool.	worker
<code>compute: platform:</code>	Required if you use compute . Use this parameter to specify the cloud provider to host the worker machines. This parameter value must match the controlPlane.platform parameter value.	aws, azure, gcp, ibmcloud, nutanix, openstack, powervs, vsphere , or <code>{}</code>
<code>compute: replicas:</code>	The number of compute machines, which are also known as worker machines, to provision.	A positive integer greater than or equal to 2 . The default value is 3 .
<code>featureSet:</code>	Enables the cluster for a feature set. A feature set is a collection of OpenShift Container Platform features that are not enabled by default. For more information about enabling a feature set during installation, see "Enabling features using feature gates".	String. The name of the feature set to enable, such as TechPreviewNoUpgrade .
<code>controlPlane:</code>	The configuration for the machines that comprise the control plane.	Array of MachinePool objects.
<code>controlPlane: architecture:</code>	Determines the instruction set architecture of the machines in the pool. Currently, clusters with varied architectures are not supported. All pools must specify the same architecture. Valid values are amd64 (the default).	String

Parameter	Description	Values
<code>controlPlane: hyperthreading:</code>	<p>Whether to enable or disable simultaneous multithreading, or hyperthreading, on control plane machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.</p>  <p>IMPORTANT</p> <p>If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.</p>	Enabled or Disabled
<code>controlPlane: name:</code>	Required if you use controlPlane . The name of the machine pool.	master
<code>controlPlane: platform:</code>	Required if you use controlPlane . Use this parameter to specify the cloud provider that hosts the control plane machines. This parameter value must match the compute.platform parameter value.	aws, azure, gcp, ibmcloud, nutanix, openstack, powervs, vsphere, or {}
<code>controlPlane: replicas:</code>	The number of control plane machines to provision.	Supported values are 3 , or 1 when deploying single-node OpenShift.
<code>credentialsMode:</code>	The Cloud Credential Operator (CCO) mode. If no mode is specified, the CCO dynamically tries to determine the capabilities of the provided credentials, with a preference for mint mode on the platforms where multiple modes are supported.	Mint, Passthrough, Manual or an empty string (<code>""</code>). ^[1]
<code>fips:</code>	Enable or disable FIPS mode. The default is false (disabled). If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default	false or true

Parameter	Description	Values
	<p data-bbox="488 107 940 210">Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.</p> <div data-bbox="488 259 592 1312" style="background-color: black; color: white; padding: 5px; margin-bottom: 10px;"> <p data-bbox="671 264 863 295">IMPORTANT</p> <p data-bbox="671 331 927 757">To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see Installing the system in FIPS mode.</p> <p data-bbox="671 792 935 1308">When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.</p> </div> <div data-bbox="488 1357 592 1559" style="background-color: black; color: white; padding: 5px;"> <p data-bbox="671 1361 762 1393">NOTE</p> <p data-bbox="671 1429 919 1554">If you are using Azure File storage, you cannot enable FIPS mode.</p> </div>	

Parameter	Description	Values
<code>imageContentSources:</code>	Sources and repositories for the release-image content.	Array of objects. Includes a source and, optionally, mirrors , as described in the following rows of this table.
<code>imageContentSources: source:</code>	Required if you use imageContentSources . Specify the repository that users refer to, for example, in image pull specifications.	String
<code>imageContentSources: mirrors:</code>	Specify one or more repositories that may also contain the same images.	Array of strings
<code>publish:</code>	How to publish or expose the user-facing endpoints of your cluster, such as the Kubernetes API, OpenShift routes.	Internal or External . To deploy a private cluster, which cannot be accessed from the internet, set publish to Internal . The default value is External .
<code>sshKey:</code>	<p>The SSH key to authenticate access to your cluster machines.</p> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p>NOTE</p> <p>For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your ssh-agent process uses.</p> </div>	For example, sshKey: ssh-ed25519 AAAA...

1. Not all CCO modes are supported for all cloud providers. For more information about CCO modes, see the "Managing cloud provider credentials" entry in the *Authentication and authorization* content.

10.10.1.4. Additional IBM Cloud configuration parameters

Additional IBM Cloud® configuration parameters are described in the following table:

Table 10.22. Additional IBM Cloud(R) parameters

Parameter	Description	Values
control Plane: platform: ibm cloud: bootVolume: encryption Key :	An IBM® Key Protect for IBM Cloud® (Key Protect) root key that should be used to encrypt the root (boot) volume of only control plane machines.	The Cloud Resource Name (CRN) of the root key. The CRN must be enclosed in quotes ("").

Parameter	Description	Values
<code>compute:platform:ibmcloud:bootVolume:encryptionKey:</code>	A Key Protect root key that should be used to encrypt the root (boot) volume of only compute machines.	The CRN of the root key. The CRN must be enclosed in quotes ("").

Parameter	Description	Values
<p>platform: ibmcloud: defaultMachinePlatform: bootvolume: encryptionKey:</p>	<p>A Key Protect root key that should be used to encrypt the root (boot) volume of all of the cluster's machines.</p> <p>When specified as part of the default machine configuration, all managed storage classes are updated with this key. As such, data volumes that are provisioned after the installation are also encrypted using this key.</p>	<p>The CRN of the root key.</p> <p>The CRN must be enclosed in quotes ("").</p>
<p>platform: ibmcloud: resourceGroupName:</p>	<p>The name of an existing resource group. By default, an installer-provisioned VPC and cluster resources are placed in this resource group. When not specified, the installation program creates the resource group for the cluster. If you are deploying the cluster into an existing VPC, the installer-provisioned cluster resources are placed in this resource group. When not specified, the installation program creates the resource group for the cluster. The VPC resources that you have provisioned must exist in a resource group that you specify using the networkResourceGroupName parameter. In either case, this resource group must only be used for a single cluster installation, as the cluster components assume ownership of all of the resources in the resource group. [1]</p>	<p>String, for example existing_resource_group.</p>

Parameter	Description	Values
platform-ibmcloud-serviceEndpoints-url	<p>A list of service endpoint names and URIs.</p> <p>By default, the installation program and cluster components use public service endpoints to access the required IBM Cloud® services.</p> <p>If network restrictions limit access to public service endpoints, you can specify an alternate service endpoint to override the default behavior.</p> <p>You can specify only one alternate service endpoint for each of the following services:</p> <ul style="list-style-type: none"> • Cloud Object Storage • DNS Services • Global Search • Global Tagging • Identity Services • Key Protect • Resource Controller • Resource Manager • VPC 	<p>A valid service endpoint name and fully qualified URI.</p> <p>Valid names include:</p> <ul style="list-style-type: none"> • COS • DNSServices • GlobalServices • GlobalTagging • IAM • KeyProtect • ResourceController • ResourceManager • VPC
platform-ibmcloud-networkResourceGroupName	<p>The name of an existing resource group. This resource contains the existing VPC and subnets to which the cluster will be deployed. This parameter is required when deploying the cluster to a VPC that you have provisioned.</p>	<p>String, for example existing_network_resource_group.</p>

Parameter	Description	Values
platform: ibmcloud: dedicatedHosts: profile:	The new dedicated host to create. If you specify a value for platform.ibmcloud.dedicatedHosts.name , this parameter is not required.	Valid IBM Cloud® dedicated host profile, such as cx2-host-152x304 . [2]
platform: ibmcloud: dedicatedHosts: name:	An existing dedicated host. If you specify a value for platform.ibmcloud.dedicatedHosts.profile , this parameter is not required.	String, for example my-dedicated-host-name .
platform: ibmcloud: type:	The instance type for all IBM Cloud® machines.	Valid IBM Cloud® instance type, such as bx2-8x32 . [2]

Parameter	Description	Values
<code>platform: ibmcloud: vpcName:</code>	The name of the existing VPC that you want to deploy your cluster to.	String.
<code>platform: ibmcloud: controlPlaneSubnets:</code>	The name(s) of the existing subnet(s) in your VPC that you want to deploy your control plane machines to. Specify a subnet for each availability zone.	String array
<code>platform: ibmcloud: computeSubnets:</code>	The name(s) of the existing subnet(s) in your VPC that you want to deploy your compute machines to. Specify a subnet for each availability zone. Subnet IDs are not supported.	String array

- Whether you define an existing resource group, or if the installer creates one, determines how the resource group is treated when the cluster is uninstalled. If you define a resource group, the installer removes all of the installer-provisioned resources, but leaves the resource group alone;

if a resource group is created as part of the installation, the installer removes all of the installer-provisioned resources and the resource group.


2. To determine which profile best meets your needs, see [Instance Profiles](#) in the IBM® documentation.

10.11. UNINSTALLING A CLUSTER ON IBM CLOUD

You can remove a cluster that you deployed to IBM Cloud®.

10.11.1. Removing a cluster that uses installer-provisioned infrastructure

You can remove a cluster that uses installer-provisioned infrastructure from your cloud.



NOTE

After uninstallation, check your cloud provider for any resources not removed properly, especially with User Provisioned Infrastructure (UPI) clusters. There might be resources that the installer did not create or that the installer is unable to access.

Prerequisites

- You have a copy of the installation program that you used to deploy the cluster.
- You have the files that the installation program generated when you created your cluster.
- You have configured the **ccoctl** binary.
- You have installed the IBM Cloud® CLI and installed or updated the VPC infrastructure service plugin. For more information see "Prerequisites" in the [IBM Cloud® CLI documentation](#).

Procedure

1. If the following conditions are met, this step is required:
 - The installer created a resource group as part of the installation process.
 - You or one of your applications created persistent volume claims (PVCs) after the cluster was deployed.

In which case, the PVCs are not removed when uninstalling the cluster, which might prevent the resource group from being successfully removed. To prevent a failure:

- a. Log in to the IBM Cloud® using the CLI.
- b. To list the PVCs, run the following command:

```
$ ibmcloud is volumes --resource-group-name <infrastructure_id>
```

For more information about listing volumes, see the [IBM Cloud® CLI documentation](#).

- c. To delete the PVCs, run the following command:

```
$ ibmcloud is volume-delete --force <volume_id>
```

For more information about deleting volumes, see the [IBM Cloud® CLI documentation](#).

- Export the API key that was created as part of the installation process.

```
$ export IC_API_KEY=<api_key>
```



NOTE

You must set the variable name exactly as specified. The installation program expects the variable name to be present to remove the service IDs that were created when the cluster was installed.

- From the directory that contains the installation program on the computer that you used to install the cluster, run the following command:

```
$ ./openshift-install destroy cluster \  
--dir <installation_directory> --log-level info 1 2
```

- For **<installation_directory>**, specify the path to the directory that you stored the installation files in.
- To view different details, specify **warn**, **debug**, or **error** instead of **info**.



NOTE

You must specify the directory that contains the cluster definition files for your cluster. The installation program requires the **metadata.json** file in this directory to delete the cluster.

- Remove the manual CCO credentials that were created for the cluster:

```
$ ccoctl ibmcloud delete-service-id \  
--credentials-requests-dir <path_to_credential_requests_directory> \  
--name <cluster_name>
```



NOTE

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

- Optional: Delete the **<installation_directory>** directory and the OpenShift Container Platform installation program.

CHAPTER 11. INSTALLING ON NUTANIX

11.1. PREPARING TO INSTALL ON NUTANIX

Before you install an OpenShift Container Platform cluster, be sure that your Nutanix environment meets the following requirements.

11.1.1. Nutanix version requirements

You must install the OpenShift Container Platform cluster to a Nutanix environment that meets the following requirements.

Table 11.1. Version requirements for Nutanix virtual environments

Component	Required version
Nutanix AOS	6.5.2.7 or later
Prism Central	pc.2022.6 or later

11.1.2. Environment requirements

Before you install an OpenShift Container Platform cluster, review the following Nutanix AOS environment requirements.

11.1.2.1. Required account privileges

The installation program requires access to a Nutanix account with the necessary permissions to deploy the cluster and to maintain the daily operation of it. The following options are available to you:

- You can use a local Prism Central user account with administrative privileges. Using a local account is the quickest way to grant access to an account with the required permissions.
- If your organization's security policies require that you use a more restrictive set of permissions, use the permissions that are listed in the following table to create a custom Cloud Native role in Prism Central. You can then assign the role to a user account that is a member of a Prism Central authentication directory.

Consider the following when managing this user account:

- When assigning entities to the role, ensure that the user can access only the Prism Element and subnet that are required to deploy the virtual machines.
- Ensure that the user is a member of the project to which it needs to assign virtual machines.

For more information, see the Nutanix documentation about creating a [Custom Cloud Native role](#), [assigning a role](#), and [adding a user to a project](#).

Example 11.1. Required permissions for creating a Custom Cloud Native role

Nutanix Object	When required	Required permissions in Nutanix API	Description
Categories	Always	Create_Category_Mapping Create_Or_Update_Name_Category Create_Or_Update_Value_Category Delete_Category_Mapping Delete_Name_Category Delete_Value_Category View_Category_Mapping View_Name_Category View_Value_Category	Create, read, and delete categories that are assigned to the OpenShift Container Platform machines.
Images	Always	Create_Image Delete_Image View_Image	Create, read, and delete the operating system images used for the OpenShift Container Platform machines.
Virtual Machines	Always	Create_Virtual_Machine Delete_Virtual_Machine View_Virtual_Machine	Create, read, and delete the OpenShift Container Platform machines.
Clusters	Always	View_Cluster	View the Prism Element clusters that host the OpenShift Container Platform machines.
Subnets	Always	View_Subnet	View the subnets that host the OpenShift Container Platform machines.

Nutanix Object	When required	Required permissions in Nutanix API	Description
Projects	If you will associate a project with compute machines, control plane machines, or all machines.	View_Project	View the projects defined in Prism Central and allow a project to be assigned to the OpenShift Container Platform machines.

11.1.2.2. Cluster limits

Available resources vary between clusters. The number of possible clusters within a Nutanix environment is limited primarily by available storage space and any limitations associated with the resources that the cluster creates, and resources that you require to deploy the cluster, such as IP addresses and networks.

11.1.2.3. Cluster resources

A minimum of 800 GB of storage is required to use a standard cluster.

When you deploy a OpenShift Container Platform cluster that uses installer-provisioned infrastructure, the installation program must be able to create several resources in your Nutanix instance. Although these resources use 856 GB of storage, the bootstrap node is destroyed as part of the installation process.

A standard OpenShift Container Platform installation creates the following resources:

- 1 label
- Virtual machines:
 - 1 disk image
 - 1 temporary bootstrap node
 - 3 control plane nodes
 - 3 compute machines

11.1.2.4. Networking requirements

You must use either AHV IP Address Management (IPAM) or Dynamic Host Configuration Protocol (DHCP) for the network and ensure that it is configured to provide persistent IP addresses to the cluster machines. Additionally, create the following networking resources before you install the OpenShift Container Platform cluster:

- IP addresses
- DNS records

**NOTE**

It is recommended that each OpenShift Container Platform node in the cluster have access to a Network Time Protocol (NTP) server that is discoverable via DHCP. Installation is possible without an NTP server. However, an NTP server prevents errors typically associated with asynchronous server clocks.

11.1.2.4.1. Required IP Addresses

An installer-provisioned installation requires two static virtual IP (VIP) addresses:

- A VIP address for the API is required. This address is used to access the cluster API.
- A VIP address for ingress is required. This address is used for cluster ingress traffic.

You specify these IP addresses when you install the OpenShift Container Platform cluster.

11.1.2.4.2. DNS records

You must create DNS records for two static IP addresses in the appropriate DNS server for the Nutanix instance that hosts your OpenShift Container Platform cluster. In each record, **<cluster_name>** is the cluster name and **<base_domain>** is the cluster base domain that you specify when you install the cluster.

If you use your own DNS or DHCP server, you must also create records for each node, including the bootstrap, control plane, and compute nodes.

A complete DNS record takes the form: **<component>.<cluster_name>.<base_domain>.**


Table 11.2. Required DNS records

Component	Record	Description
API VIP	api.<cluster_name>.<base_domain>.	This DNS A/AAAA or CNAME record must point to the load balancer for the control plane machines. This record must be resolvable by both clients external to the cluster and from all the nodes within the cluster.
Ingress VIP	*.apps.<cluster_name>.<base_domain>.	A wildcard DNS A/AAAA or CNAME record that points to the load balancer that targets the machines that run the Ingress router pods, which are the worker nodes by default. This record must be resolvable by both clients external to the cluster and from all the nodes within the cluster.

11.1.3. Configuring the Cloud Credential Operator utility

The Cloud Credential Operator (CCO) manages cloud provider credentials as Kubernetes custom resource definitions (CRDs). To install a cluster on Nutanix, you must set the CCO to **manual** mode as part of the installation process.

To create and manage cloud credentials from outside of the cluster when the Cloud Credential Operator (CCO) is operating in manual mode, extract and prepare the CCO utility (**ccoctl**) binary.



NOTE

The **ccoctl** utility is a Linux binary that must run in a Linux environment.

Prerequisites

- You have access to an OpenShift Container Platform account with cluster administrator access.
- You have installed the OpenShift CLI (**oc**).


Procedure

1. Set a variable for the OpenShift Container Platform release image by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

2. Obtain the CCO container image from the OpenShift Container Platform release image by running the following command:

```
$ CCO_IMAGE=$(oc adm release info --image-for='cloud-credential-operator'
$RELEASE_IMAGE -a ~/.pull-secret)
```



NOTE

Ensure that the architecture of the **\$RELEASE_IMAGE** matches the architecture of the environment in which you will use the **ccoctl** tool.

3. Extract the **ccoctl** binary from the CCO container image within the OpenShift Container Platform release image by running the following command:

```
$ oc image extract $CCO_IMAGE \
--file="/usr/bin/ccoctl.<rhel_version>" \
-a ~/.pull-secret
```

- 1 For **<rhel_version>**, specify the value that corresponds to the version of Red Hat Enterprise Linux (RHEL) that the host uses. If no value is specified, **ccoctl.rhel8** is used by default. The following values are valid:

- **rhel8**: Specify this value for hosts that use RHEL 8.
- **rhel9**: Specify this value for hosts that use RHEL 9.

4. Change the permissions to make **ccoctl** executable by running the following command:

```
$ chmod 775 ccoctl.<rhel_version>
```

Verification

- To verify that **ccoctl** is ready to use, display the help file by running the following command:

```
$ ccoctl --help
```

Output of **ccoctl --help**

```
OpenShift credentials provisioning tool
```

```
Usage:
```

```
ccoctl [command]
```

```
Available Commands:
```

```
aws      Manage credentials objects for AWS cloud
azure    Manage credentials objects for Azure
gcp      Manage credentials objects for Google cloud
help     Help about any command
ibmcloud Manage credentials objects for IBM Cloud
nutanix  Manage credentials objects for Nutanix
```

```
Flags:
```

```
-h, --help  help for ccoctl
```

```
Use "ccoctl [command] --help" for more information about a command.
```

Additional resources

- [Preparing to update a cluster with manually maintained credentials](#)

11.2. FAULT TOLERANT DEPLOYMENTS USING MULTIPLE PRISM ELEMENTS

By default, the installation program installs control plane and compute machines into a single Nutanix Prism Element (cluster). To improve the fault tolerance of your OpenShift Container Platform cluster, you can specify that these machines be distributed across multiple Nutanix clusters by configuring failure domains.

A failure domain represents an additional Prism Element instance that is available to OpenShift Container Platform machine pools during and after installation.

11.2.1. Failure domain requirements

When planning to use failure domains, consider the following requirements:

- All Nutanix Prism Element instances must be managed by the same instance of Prism Central. A deployment that is comprised of multiple Prism Central instances is not supported.
- The machines that make up the Prism Element clusters must reside on the same Ethernet network for failure domains to be able to communicate with each other.

- A subnet is required in each Prism Element that will be used as a failure domain in the OpenShift Container Platform cluster. When defining these subnets, they must share the same IP address prefix (CIDR) and should contain the virtual IP addresses that the OpenShift Container Platform cluster uses.

11.2.2. Installation method and failure domain configuration

The OpenShift Container Platform installation method determines how and when you configure failure domains:

- If you deploy using installer-provisioned infrastructure, you can configure failure domains in the installation configuration file before deploying the cluster. For more information, see [Configuring failure domains](#).
You can also configure failure domains after the cluster is deployed. For more information about configuring failure domains post-installation, see [Adding failure domains to an existing Nutanix cluster](#).
- If you deploy using infrastructure that you manage (user-provisioned infrastructure) no additional configuration is required. After the cluster is deployed, you can manually distribute control plane and compute machines across failure domains.

11.3. INSTALLING A CLUSTER ON NUTANIX

In OpenShift Container Platform version 4.16, you can choose one of the following options to install a cluster on your Nutanix instance:

Using installer-provisioned infrastructure: Use the procedures in the following sections to use installer-provisioned infrastructure. Installer-provisioned infrastructure is ideal for installing in connected or disconnected network environments. The installer-provisioned infrastructure includes an installation program that provisions the underlying infrastructure for the cluster.

Using the Assisted Installer: The [Assisted Installer](#) hosted at console.redhat.com. The Assisted Installer cannot be used in disconnected environments. The Assisted Installer does not provision the underlying infrastructure for the cluster, so you must provision the infrastructure before the running the Assisted Installer. Installing with the Assisted Installer also provides integration with Nutanix, enabling autoscaling. See [Installing an on-premise cluster using the Assisted Installer](#) for additional details.

Using user-provisioned infrastructure: Complete the relevant steps outlined in the [Installing a cluster on any platform](#) documentation.

11.3.1. Prerequisites

- You have reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- The installation program requires access to port 9440 on Prism Central and Prism Element. You verified that port 9440 is accessible.
- If you use a firewall, you have met these prerequisites:
 - You confirmed that port 9440 is accessible. Control plane nodes must be able to reach Prism Central and Prism Element on port 9440 for the installation to succeed.
 - You configured the firewall to [grant access](#) to the sites that OpenShift Container Platform requires. This includes the use of Telemetry.

- If your Nutanix environment is using the default self-signed SSL certificate, replace it with a certificate that is signed by a CA. The installation program requires a valid CA-signed certificate to access to the Prism Central API. For more information about replacing the self-signed certificate, see the [Nutanix AOS Security Guide](#).

If your Nutanix environment uses an internal CA to issue certificates, you must configure a cluster-wide proxy as part of the installation process. For more information, see [Configuring a custom PKI](#).



IMPORTANT

Use 2048-bit certificates. The installation fails if you use 4096-bit certificates with Prism Central 2022.x.

11.3.2. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.



IMPORTANT

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

11.3.3. Internet access for Prism Central

Prism Central requires internet access to obtain the Red Hat Enterprise Linux CoreOS (RHCOS) image that is required to install the cluster. The RHCOS image for Nutanix is available at rhcos.mirror.openshift.com.

11.3.4. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.



NOTE

You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.



NOTE

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.



NOTE

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```



NOTE

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

11.3.5. Obtaining the installation program


Before you install OpenShift Container Platform, download the installation file on the host you are using for installation.

Prerequisites


- You have a computer that runs Linux or macOS, with 500 MB of local disk space.

Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select your infrastructure provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.

**IMPORTANT**

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.

**IMPORTANT**

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.

4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#) . This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

11.3.6. Adding Nutanix root CA certificates to your system trust

Because the installation program requires access to the Prism Central API, you must add your Nutanix trusted root CA certificates to your system trust before you install an OpenShift Container Platform cluster.

Procedure

1. From the Prism Central web console, download the Nutanix root CA certificates.
2. Extract the compressed file that contains the Nutanix root CA certificates.
3. Add the files for your operating system to the system trust. For example, on a Fedora operating system, run the following command:

```
# cp certs/lin/* /etc/pki/ca-trust/source/anchors
```

4. Update your system trust. For example, on a Fedora operating system, run the following command:

```
# update-ca-trust extract
```

11.3.7. Creating the installation configuration file

You can customize the OpenShift Container Platform cluster you install on Nutanix.

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have verified that you have met the Nutanix networking requirements. For more information, see "Preparing to install on Nutanix".

Procedure

1. Create the **install-config.yaml** file.
 - a. Change to the directory that contains the installation program and run the following command:

```
┆ $ ./openshift-install create install-config --dir <installation_directory> 1
```

- 1 For **<installation_directory>**, specify the directory name to store the files that the installation program creates.

When specifying the directory:

- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
- Use an empty directory. Some installation assets, such as bootstrap X.509 certificates, have short expiration intervals, therefore you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

- b. At the prompts, provide the configuration details for your cloud:
 - i. Optional: Select an SSH key to use to access your cluster machines.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.


- ii. Select **nutanix** as the platform to target.
- iii. Enter the Prism Central domain name or IP address.
- iv. Enter the port that is used to log into Prism Central.
- v. Enter the credentials that are used to log into Prism Central.
The installation program connects to Prism Central.
- vi. Select the Prism Element that will manage the OpenShift Container Platform cluster.
- vii. Select the network subnet to use.
- viii. Enter the virtual IP address that you configured for control plane API access.

- ix. Enter the virtual IP address that you configured for cluster ingress.
 - x. Enter the base domain. This base domain must be the same one that you configured in the DNS records.
 - xi. Enter a descriptive name for your cluster.
The cluster name you enter must match the cluster name you specified when configuring the DNS records.
2. Optional: Update one or more of the default configuration parameters in the **install.config.yaml** file to customize the installation.
For more information about the parameters, see "Installation configuration parameters".

**NOTE**

If you are installing a three-node cluster, be sure to set the **compute.replicas** parameter to **0**. This ensures that cluster's control planes are schedulable. For more information, see "Installing a three-node cluster on Nutanix".

3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.

**IMPORTANT**


The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

Additional resources

- [Installation configuration parameters for Nutanix](#)

11.3.7.1. Sample customized install-config.yaml file for Nutanix

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.

**IMPORTANT**

This sample YAML file is provided for reference only. You must obtain your **install-config.yaml** file by using the installation program and modify it.

```

apiVersion: v1
baseDomain: example.com 1
compute: 2
- hyperthreading: Enabled 3
  name: worker
  replicas: 3
platform:
  nutanix: 4
    cpus: 2
    coresPerSocket: 2
    memoryMiB: 8196
    osDisk:
      diskSizeGiB: 120

```

```
categories: 5
  - key: <category_key_name>
    value: <category_value>
controlPlane: 6
hyperthreading: Enabled 7
name: master
replicas: 3
platform:
  nutanix: 8
    cpus: 4
    coresPerSocket: 2
    memoryMiB: 16384
    osDisk:
      diskSizeGiB: 120
    categories: 9
      - key: <category_key_name>
        value: <category_value>
metadata:
  creationTimestamp: null
  name: test-cluster 10
networking:
  clusterNetwork:
    - cidr: 10.128.0.0/14
      hostPrefix: 23
  machineNetwork:
    - cidr: 10.0.0.0/16
  networkType: OVNKubernetes 11
  serviceNetwork:
    - 172.30.0.0/16
platform:
  nutanix:
    apiVIPs:
      - 10.40.142.7 12
    defaultMachinePlatform:
      bootType: Legacy
      categories: 13
        - key: <category_key_name>
          value: <category_value>
      project: 14
        type: name
        name: <project_name>
    ingressVIPs:
      - 10.40.142.8 15
    prismCentral:
      endpoint:
        address: your.prismcentral.domainname 16
        port: 9440 17
      password: <password> 18
      username: <username> 19
    prismElements:
      - endpoint:
          address: your.prismelement.domainname
          port: 9440
          uuid: 0005b0f1-8f43-a0f2-02b7-3cecef193712
```

```

subnetUUIDs:
- c7938dc6-7659-453e-a688-e26020c68e43
clusterOSImage: http://example.com/images/rhcos-47.83.202103221318-0-nutanix.x86_64.qcow2

```

20

credentialsMode: Manual

publish: External

pullSecret: '{"auths": ...}' **21**fips: false **22**sshKey: ssh-ed25519 AAAA... **23**

1 10 12 15 16 17 18 19 21 Required. The installation program prompts you for this value.

2 6 The **controlPlane** section is a single mapping, but the compute section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Although both sections currently define a single machine pool, it is possible that future versions of OpenShift Container Platform will support defining multiple compute pools during installation. Only one control plane pool is used.

3 7 Whether to enable or disable simultaneous multithreading, or **hyperthreading**. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores. You can disable it by setting the parameter value to **Disabled**. If you disable simultaneous multithreading in some cluster machines, you must disable it in all cluster machines.



IMPORTANT

If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.

4 8 Optional: Provide additional configuration for the machine pool parameters for the compute and control plane machines.

5 9 13 Optional: Provide one or more pairs of a prism category key and a prism category value. These category key-value pairs must exist in Prism Central. You can provide separate categories to compute machines, control plane machines, or all machines.

11 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.

14 Optional: Specify a project with which VMs are associated. Specify either **name** or **uuid** for the project type, and then provide the corresponding UUID or project name. You can associate projects to compute machines, control plane machines, or all machines.

20 Optional: By default, the installation program downloads and installs the Red Hat Enterprise Linux CoreOS (RHCOS) image. If Prism Central does not have internet access, you can override the default behavior by hosting the RHCOS image on any HTTP server and pointing the installation program to the image.

22 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.

**IMPORTANT**

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

23

Optional: You can provide the **sshKey** value that you use to access the machines in your cluster.

**NOTE**

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

11.3.7.2. Configuring failure domains

Failure domains improve the fault tolerance of an OpenShift Container Platform cluster by distributing control plane and compute machines across multiple Nutanix Prism Elements (clusters).

TIP

It is recommended that you configure three failure domains to ensure high-availability.

Prerequisites

- You have an installation configuration file (**install-config.yaml**).

Procedure

1. Edit the **install-config.yaml** file and add the following stanza to configure the first failure domain:

```

apiVersion: v1
baseDomain: example.com
compute:
# ...
platform:
  nutanix:
    failureDomains:
      - name: <failure_domain_name>
        prismElement:
          name: <prism_element_name>
          uuid: <prism_element_uuid>
        subnetUUIDs:
          - <network_uuid>
# ...

```

where:

<failure_domain_name>

Specifies a unique name for the failure domain. The name is limited to 64 or fewer characters, which can include lower-case letters, digits, and a dash (-). The dash cannot be in the leading or ending position of the name.

<prism_element_name>

Optional. Specifies the name of the Prism Element.

<prism_element_uuid>

Specifies the UUID of the Prism Element.

<network_uuid>

Specifies the UUID of the Prism Element subnet object. The subnet's IP address prefix (CIDR) should contain the virtual IP addresses that the OpenShift Container Platform cluster uses. Only one subnet per failure domain (Prism Element) in an OpenShift Container Platform cluster is supported.

2. As required, configure additional failure domains.
3. To distribute control plane and compute machines across the failure domains, do one of the following:
 - If compute and control plane machines can share the same set of failure domains, add the failure domain names under the cluster's default machine configuration.

Example of control plane and compute machines sharing a set of failure domains

```
apiVersion: v1
baseDomain: example.com
compute:
# ...
platform:
  nutanix:
    defaultMachinePlatform:
      failureDomains:
        - failure-domain-1
        - failure-domain-2
        - failure-domain-3
# ...
```

- If compute and control plane machines must use different failure domains, add the failure domain names under the respective machine pools.

Example of control plane and compute machines using different failure domains

```
apiVersion: v1
baseDomain: example.com
compute:
# ...
controlPlane:
  platform:
    nutanix:
      failureDomains:
        - failure-domain-1
        - failure-domain-2
        - failure-domain-3
# ...
compute:
```

```
platform:
  nutanix:
    failureDomains:
      - failure-domain-1
      - failure-domain-2
# ...
```

4. Save the file.

11.3.7.3. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.

- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with `.` to match subdomains only. For example, `.y.com` matches `x.y.com`, but not `y.com`. Use `*` to bypass the proxy for all destinations.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.

**NOTE**

The installation program does not support the proxy **readinessEndpoints** field.

**NOTE**

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.

**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

11.3.8. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.

**IMPORTANT**

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**.
To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.



NOTE

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your PATH.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

11.3.9. Configuring IAM for Nutanix

Installing the cluster requires that the Cloud Credential Operator (CCO) operate in manual mode. While the installation program configures the CCO for manual mode, you must specify the identity and access management secrets.

Prerequisites

- You have configured the **ccoctl** binary.
- You have an **install-config.yaml** file.

Procedure

1. Create a YAML file that contains the credentials data in the following format:

Credentials data format

```
credentials:
- type: basic_auth 1
  data:
    prismCentral: 2
      username: <username_for_prism_central>
      password: <password_for_prism_central>
```

```
prismElements: 3
  - name: <name_of_prism_element>
    username: <username_for_prism_element>
    password: <password_for_prism_element>
```

- 1** Specify the authentication type. Only basic authentication is supported.
 - 2** Specify the Prism Central credentials.
 - 3** Optional: Specify the Prism Element credentials.
2. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

3. Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included 1 \
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml 2 \
  --to=<path_to_directory_for_credentials_requests> 3
```

- 1** The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2** Specify the location of the **install-config.yaml** file.
- 3** Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

Sample **CredentialsRequest** object

```
apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  annotations:
    include.release.openshift.io/self-managed-high-availability: "true"
  labels:
    controller-tools.k8s.io: "1.0"
  name: openshift-machine-api-nutanix
  namespace: openshift-cloud-credential-operator
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: NutanixProviderSpec
  secretRef:
    name: nutanix-credentials
    namespace: openshift-machine-api
```

4. Use the **ccoctl** tool to process all **CredentialsRequest** objects by running the following command:

```
$ ccoctl nutanix create-shared-secrets \
  --credentials-requests-dir=<path_to_credentials_requests_directory> \ 1
  --output-dir=<ccoctl_output_dir> \ 2
  --credentials-source-filepath=<path_to_credentials_file> \ 3
```

- 1** Specify the path to the directory that contains the files for the component **CredentialsRequests** objects.
- 2** Optional: Specify the directory in which you want the **ccoctl** utility to create objects. By default, the utility creates objects in the directory in which the commands are run.
- 3** Optional: Specify the directory that contains the credentials data YAML file. By default, **ccoctl** expects this file to be in **<home_directory>/./nutanix/credentials**.

5. Edit the **install-config.yaml** configuration file so that the **credentialsMode** parameter is set to **Manual**.

Example install-config.yaml configuration file

```
apiVersion: v1
baseDomain: cluster1.example.com
credentialsMode: Manual 1
...
```

- 1** Add this line to set the **credentialsMode** parameter to **Manual**.

6. Create the installation manifests by running the following command:

```
$ openshift-install create manifests --dir <installation_directory> 1
```

- 1** Specify the path to the directory that contains the **install-config.yaml** file for your cluster.

7. Copy the generated credential files to the target manifests directory by running the following command:

```
$ cp <ccoctl_output_dir>/manifests/*credentials.yaml ./<installation_directory>/manifests
```

Verification

- Ensure that the appropriate secrets exist in the **manifests** directory.

```
$ ls ./<installation_directory>/manifests
```

Example output

```
cluster-config.yaml
cluster-dns-02-config.yml
```

```

cluster-infrastructure-02-config.yml
cluster-ingress-02-config.yml
cluster-network-01-crd.yml
cluster-network-02-config.yml
cluster-proxy-01-config.yml
cluster-scheduler-02-config.yml
cvo-overrides.yml
kube-cloud-config.yml
kube-system-configmap-root-ca.yml
machine-config-server-tls-secret.yml
openshift-config-secret-pull-secret.yml
openshift-cloud-controller-manager-nutanix-credentials-credentials.yml
openshift-machine-api-nutanix-credentials-credentials.yml

```

11.3.10. Adding config map and secret resources required for Nutanix CCM

Installations on Nutanix require additional **ConfigMap** and **Secret** resources to integrate with the Nutanix Cloud Controller Manager (CCM).

Prerequisites

- You have created a **manifests** directory within your installation directory.

Procedure

- Navigate to the **manifests** directory:

```
$ cd <path_to_installation_directory>/manifests
```

- Create the **cloud-conf ConfigMap** file with the name **openshift-cloud-controller-manager-cloud-config.yml** and add the following information:

```

apiVersion: v1
kind: ConfigMap
metadata:
  name: cloud-conf
  namespace: openshift-cloud-controller-manager
data:
  cloud.conf: "{
    \"prismCentral\": {
      \"address\": \"<prism_central_FQDN/IP>\", 1
      \"port\": 9440,
      \"credentialRef\": {
        \"kind\": \"Secret\",
        \"name\": \"nutanix-credentials\",
        \"namespace\": \"openshift-cloud-controller-manager\"
      }
    },
    \"topologyDiscovery\": {
      \"type\": \"Prism\",
      \"topologyCategories\": null
    },
    \"enableCustomLabeling\": true
  }"

```



1 Specify the Prism Central FQDN/IP.

3. Verify that the file **cluster-infrastructure-02-config.yml** exists and has the following information:

```
spec:
  cloudConfig:
    key: config
    name: cloud-provider-config
```

11.3.11. Services for a user-managed load balancer

You can configure an OpenShift Container Platform cluster to use a user-managed load balancer in place of the default load balancer.



IMPORTANT

Configuring a user-managed load balancer depends on your vendor's load balancer.


The information and examples in this section are for guideline purposes only. Consult the vendor documentation for more specific information about the vendor's load balancer.

Red Hat supports the following services for a user-managed load balancer:

- Ingress Controller
- OpenShift API
- OpenShift MachineConfig API

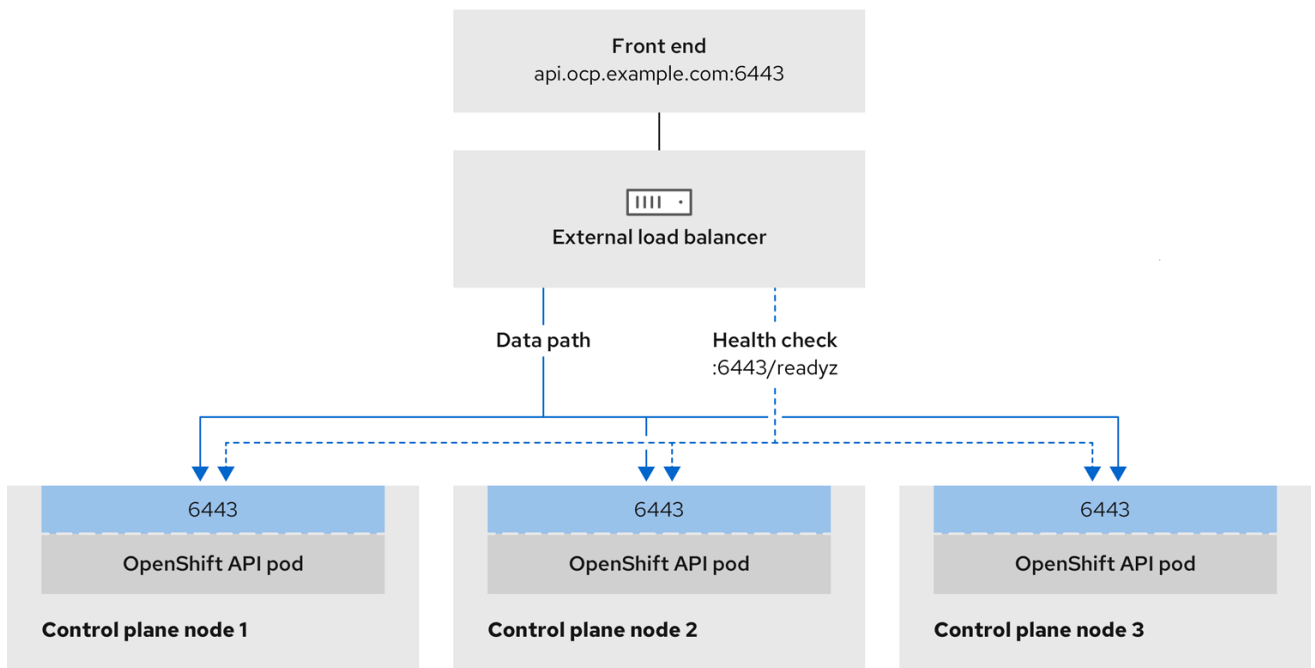
You can choose whether you want to configure one or all of these services for a user-managed load balancer. Configuring only the Ingress Controller service is a common configuration option. To better understand each service, view the following diagrams:

Figure 11.1. Example network workflow that shows an Ingress Controller operating in an OpenShift Container Platform environment




496_OpenShift_1223

Figure 11.2. Example network workflow that shows an OpenShift API operating in an OpenShift Container Platform environment



496_OpenShift_1223

Figure 11.3. Example network workflow that shows an OpenShift MachineConfig API operating in an OpenShift Container Platform environment



The following configuration options are supported for user-managed load balancers:

- Use a node selector to map the Ingress Controller to a specific set of nodes. You must assign a static IP address to each node in this set, or configure each node to receive the same IP address from the Dynamic Host Configuration Protocol (DHCP). Infrastructure nodes commonly receive this type of configuration.
- Target all IP addresses on a subnet. This configuration can reduce maintenance overhead, because you can create and destroy nodes within those networks without reconfiguring the load balancer targets. If you deploy your ingress pods by using a machine set on a smaller network, such as a `/27` or `/28`, you can simplify your load balancer targets.

TIP

You can list all IP addresses that exist in a network by checking the machine config pool's resources.


Before you configure a user-managed load balancer for your OpenShift Container Platform cluster, consider the following information:

- For a front-end IP address, you can use the same IP address for the front-end IP address, the Ingress Controller's load balancer, and API load balancer. Check the vendor's documentation for this capability.
- For a back-end IP address, ensure that an IP address for an OpenShift Container Platform control plane node does not change during the lifetime of the user-managed load balancer. You can achieve this by completing one of the following actions:
 - Assign a static IP address to each control plane node.

- Configure each node to receive the same IP address from the DHCP every time the node requests a DHCP lease. Depending on the vendor, the DHCP lease might be in the form of an IP reservation or a static DHCP assignment.
- Manually define each node that runs the Ingress Controller in the user-managed load balancer for the Ingress Controller back-end service. For example, if the Ingress Controller moves to an undefined node, a connection outage can occur.

11.3.11.1. Configuring a user-managed load balancer

You can configure an OpenShift Container Platform cluster to use a user-managed load balancer in place of the default load balancer.



IMPORTANT

Before you configure a user-managed load balancer, ensure that you read the "Services for a user-managed load balancer" section.

Read the following prerequisites that apply to the service that you want to configure for your user-managed load balancer.



NOTE

MetaLB, which runs on a cluster, functions as a user-managed load balancer.

OpenShift API prerequisites

- You defined a front-end IP address.
- TCP ports 6443 and 22623 are exposed on the front-end IP address of your load balancer. Check the following items:
 - Port 6443 provides access to the OpenShift API service.
 - Port 22623 can provide ignition startup configurations to nodes.
- The front-end IP address and port 6443 are reachable by all users of your system with a location external to your OpenShift Container Platform cluster.
- The front-end IP address and port 22623 are reachable only by OpenShift Container Platform nodes.
- The load balancer backend can communicate with OpenShift Container Platform control plane nodes on port 6443 and 22623.

Ingress Controller prerequisites

- You defined a front-end IP address.
- TCP ports 443 and 80 are exposed on the front-end IP address of your load balancer.
- The front-end IP address, port 80 and port 443 are be reachable by all users of your system with a location external to your OpenShift Container Platform cluster.

- The front-end IP address, port 80 and port 443 are reachable to all nodes that operate in your OpenShift Container Platform cluster.
- The load balancer backend can communicate with OpenShift Container Platform nodes that run the Ingress Controller on ports 80, 443, and 1936.

Prerequisite for health check URL specifications

You can configure most load balancers by setting health check URLs that determine if a service is available or unavailable. OpenShift Container Platform provides these health checks for the OpenShift API, Machine Configuration API, and Ingress Controller backend services.

The following examples show health check specifications for the previously listed backend services:

Example of a Kubernetes API health check specification

```
Path: HTTPS:6443/readyz
Healthy threshold: 2
Unhealthy threshold: 2
Timeout: 10
Interval: 10
```

Example of a Machine Config API health check specification

```
Path: HTTPS:22623/healthz
Healthy threshold: 2
Unhealthy threshold: 2
Timeout: 10
Interval: 10
```

Example of an Ingress Controller health check specification

```
Path: HTTP:1936/healthz/ready
Healthy threshold: 2
Unhealthy threshold: 2
Timeout: 5
Interval: 10
```

Procedure

1. Configure the HAProxy Ingress Controller, so that you can enable access to the cluster from your load balancer on ports 6443, 22623, 443, and 80. Depending on your needs, you can specify the IP address of a single subnet or IP addresses from multiple subnets in your HAProxy configuration.

Example HAProxy configuration with one listed subnet

```
# ...
listen my-cluster-api-6443
  bind 192.168.1.100:6443
  mode tcp
  balance roundrobin
  option httpchk
  http-check connect
  http-check send meth GET uri /readyz
```

```

http-check expect status 200
  server my-cluster-master-2 192.168.1.101:6443 check inter 10s rise 2 fall 2
  server my-cluster-master-0 192.168.1.102:6443 check inter 10s rise 2 fall 2
  server my-cluster-master-1 192.168.1.103:6443 check inter 10s rise 2 fall 2

listen my-cluster-machine-config-api-22623
  bind 192.168.1.100:22623
  mode tcp
  balance roundrobin
  option httpchk
  http-check connect
  http-check send meth GET uri /healthz
  http-check expect status 200
    server my-cluster-master-2 192.168.1.101:22623 check inter 10s rise 2 fall 2
    server my-cluster-master-0 192.168.1.102:22623 check inter 10s rise 2 fall 2
    server my-cluster-master-1 192.168.1.103:22623 check inter 10s rise 2 fall 2

listen my-cluster-apps-443
  bind 192.168.1.100:443
  mode tcp
  balance roundrobin
  option httpchk
  http-check connect
  http-check send meth GET uri /healthz/ready
  http-check expect status 200
    server my-cluster-worker-0 192.168.1.111:443 check port 1936 inter 10s rise 2 fall 2
    server my-cluster-worker-1 192.168.1.112:443 check port 1936 inter 10s rise 2 fall 2
    server my-cluster-worker-2 192.168.1.113:443 check port 1936 inter 10s rise 2 fall 2

listen my-cluster-apps-80
  bind 192.168.1.100:80
  mode tcp
  balance roundrobin
  option httpchk
  http-check connect
  http-check send meth GET uri /healthz/ready
  http-check expect status 200
    server my-cluster-worker-0 192.168.1.111:80 check port 1936 inter 10s rise 2 fall 2
    server my-cluster-worker-1 192.168.1.112:80 check port 1936 inter 10s rise 2 fall 2
    server my-cluster-worker-2 192.168.1.113:80 check port 1936 inter 10s rise 2 fall 2
# ...

```

Example HAProxy configuration with multiple listed subnets

```

# ...
listen api-server-6443
  bind *:6443
  mode tcp
    server master-00 192.168.83.89:6443 check inter 1s
    server master-01 192.168.84.90:6443 check inter 1s
    server master-02 192.168.85.99:6443 check inter 1s
    server bootstrap 192.168.80.89:6443 check inter 1s

listen machine-config-server-22623
  bind *:22623
  mode tcp

```

```

server master-00 192.168.83.89:22623 check inter 1s
server master-01 192.168.84.90:22623 check inter 1s
server master-02 192.168.85.99:22623 check inter 1s
server bootstrap 192.168.80.89:22623 check inter 1s

listen ingress-router-80
  bind *:80
  mode tcp
  balance source
    server worker-00 192.168.83.100:80 check inter 1s
    server worker-01 192.168.83.101:80 check inter 1s

listen ingress-router-443
  bind *:443
  mode tcp
  balance source
    server worker-00 192.168.83.100:443 check inter 1s
    server worker-01 192.168.83.101:443 check inter 1s

listen ironic-api-6385
  bind *:6385
  mode tcp
  balance source
    server master-00 192.168.83.89:6385 check inter 1s
    server master-01 192.168.84.90:6385 check inter 1s
    server master-02 192.168.85.99:6385 check inter 1s
    server bootstrap 192.168.80.89:6385 check inter 1s

listen inspector-api-5050
  bind *:5050
  mode tcp
  balance source
    server master-00 192.168.83.89:5050 check inter 1s
    server master-01 192.168.84.90:5050 check inter 1s
    server master-02 192.168.85.99:5050 check inter 1s
    server bootstrap 192.168.80.89:5050 check inter 1s
# ...

```

2. Use the **curl** CLI command to verify that the user-managed load balancer and its resources are operational:
 - a. Verify that the cluster machine configuration API is accessible to the Kubernetes API server resource, by running the following command and observing the response:

```
$ curl https://<loadbalancer_ip_address>:6443/version --insecure
```

If the configuration is correct, you receive a JSON object in response:

```
{
  "major": "1",
  "minor": "11+",
  "gitVersion": "v1.11.0+ad103ed",
  "gitCommit": "ad103ed",
  "gitTreeState": "clean",
  "buildDate": "2019-01-09T06:44:10Z",
  "goVersion": "go1.10.3",

```

```
"compiler": "gc",
"platform": "linux/amd64"
}
```

- b. Verify that the cluster machine configuration API is accessible to the Machine config server resource, by running the following command and observing the output:

```
$ curl -v https://<loadbalancer_ip_address>:22623/healthz --insecure
```

If the configuration is correct, the output from the command shows the following response:

```
HTTP/1.1 200 OK
Content-Length: 0
```

- c. Verify that the controller is accessible to the Ingress Controller resource on port 80, by running the following command and observing the output:

```
$ curl -I -L -H "Host: console-openshift-console.apps.<cluster_name>.<base_domain>"
http://<load_balancer_front_end_IP_address>
```

If the configuration is correct, the output from the command shows the following response:

```
HTTP/1.1 302 Found
content-length: 0
location: https://console-openshift-console.apps.ocp4.private.opequon.net/
cache-control: no-cache
```

- d. Verify that the controller is accessible to the Ingress Controller resource on port 443, by running the following command and observing the output:

```
$ curl -I -L --insecure --resolve console-openshift-console.apps.<cluster_name>.<base_domain>:443:<Load Balancer Front End IP Address> https://console-openshift-console.apps.<cluster_name>.<base_domain>
```

If the configuration is correct, the output from the command shows the following response:


```
HTTP/1.1 200 OK
referrer-policy: strict-origin-when-cross-origin
set-cookie: csrf-
token=UIYWOyQ62LWjw2h003xtYSKlh1a0Py2hhctw0WmV2YEdhJfFyQwWcGBsja261dGLgaYO0nxzVERhiXt6QepA7g==; Path=/; Secure; SameSite=Lax
x-content-type-options: nosniff
x-dns-prefetch-control: off
x-frame-options: DENY
x-xss-protection: 1; mode=block
date: Wed, 04 Oct 2023 16:29:38 GMT
content-type: text/html; charset=utf-8
set-cookie:
1e2670d92730b515ce3a1bb65da45062=1bf5e9573c9a2760c964ed1659cc1673; path=/;
HttpOnly; Secure; SameSite=None
cache-control: private
```


- Configure the DNS records for your cluster to target the front-end IP addresses of the user-managed load balancer. You must update records to your DNS server for the cluster API and applications over the load balancer.

Examples of modified DNS records

```
<load_balancer_ip_address> A api.<cluster_name>.<base_domain>
A record pointing to Load Balancer Front End
```

```
<load_balancer_ip_address> A apps.<cluster_name>.<base_domain>
A record pointing to Load Balancer Front End
```



IMPORTANT

DNS propagation might take some time for each DNS record to become available. Ensure that each DNS record propagates before validating each record.

- For your OpenShift Container Platform cluster to use the user-managed load balancer, you must specify the following configuration in your cluster's **install-config.yaml** file:

```
# ...
platform:
  nutanix:
    loadBalancer:
      type: UserManaged 1
      apiVIPs:
        - <api_ip> 2
      ingressVIPs:
        - <ingress_ip> 3
# ...
```

- Set **UserManaged** for the **type** parameter to specify a user-managed load balancer for your cluster. The parameter defaults to **OpenShiftManagedDefault**, which denotes the default internal load balancer. For services defined in an **openshift-kni-infra** namespace, a user-managed load balancer can deploy the **coredns** service to pods in your cluster but ignores **keepalived** and **haproxy** services.
- Required parameter when you specify a user-managed load balancer. Specify the user-managed load balancer's public IP address, so that the Kubernetes API can communicate with the user-managed load balancer.
- Required parameter when you specify a user-managed load balancer. Specify the user-managed load balancer's public IP address, so that the user-managed load balancer can manage ingress traffic for your cluster.

Verification

- Use the **curl** CLI command to verify that the user-managed load balancer and DNS record configuration are operational:
 - Verify that you can access the cluster API, by running the following command and observing the output:

```
$ curl https://api.<cluster_name>.<base_domain>:6443/version --insecure
```

If the configuration is correct, you receive a JSON object in response:

```
{
  "major": "1",
  "minor": "11+",
  "gitVersion": "v1.11.0+ad103ed",
  "gitCommit": "ad103ed",
  "gitTreeState": "clean",
  "buildDate": "2019-01-09T06:44:10Z",
  "goVersion": "go1.10.3",
  "compiler": "gc",
  "platform": "linux/amd64"
}
```

- b. Verify that you can access the cluster machine configuration, by running the following command and observing the output:

```
$ curl -v https://api.<cluster_name>.<base_domain>:22623/healthz --insecure
```

If the configuration is correct, the output from the command shows the following response:

```
HTTP/1.1 200 OK
Content-Length: 0
```

- c. Verify that you can access each cluster application on port, by running the following command and observing the output:

```
$ curl http://console-openshift-console.apps.<cluster_name>.<base_domain> -I -L --insecure
```

If the configuration is correct, the output from the command shows the following response:

```
HTTP/1.1 302 Found
content-length: 0
location: https://console-openshift-console.apps.<cluster-name>.<base domain>/
cache-control: no-cacheHTTP/1.1 200 OK
referrer-policy: strict-origin-when-cross-origin
set-cookie: csrf-
token=39HoZgztDnzjJkq/JuLJMeoKNXIfiVv2YgZc09c3TBOBU4NI6kDXaJH1LdicNhN1UsQ
Wzon4Dor9GWGfopaTEQ==; Path=/; Secure
x-content-type-options: nosniff
x-dns-prefetch-control: off
x-frame-options: DENY
x-xss-protection: 1; mode=block
date: Tue, 17 Nov 2020 08:42:10 GMT
content-type: text/html; charset=utf-8
set-cookie:
1e2670d92730b515ce3a1bb65da45062=9b714eb87e93cf34853e87a92d6894be; path=/;
HttpOnly; Secure; SameSite=None
cache-control: private
```

- d. Verify that you can access each cluster application on port 443, by running the following command and observing the output:

```
$ curl https://console-openshift-console.apps.<cluster_name>.<base_domain> -I -L --insecure
```

If the configuration is correct, the output from the command shows the following response:

```
HTTP/1.1 200 OK
referrer-policy: strict-origin-when-cross-origin
set-cookie: csrf-
token=UIYW0yQ62LWjw2h003xtYSKlh1a0Py2hhctw0WmV2YEdhJfYqWwCGBsja261dG
LgaYO0nxzVERhiXt6QepA7g==; Path=/; Secure; SameSite=Lax
x-content-type-options: nosniff
x-dns-prefetch-control: off
x-frame-options: DENY
x-xss-protection: 1; mode=block
date: Wed, 04 Oct 2023 16:29:38 GMT
content-type: text/html; charset=utf-8
set-cookie:
1e2670d92730b515ce3a1bb65da45062=1bf5e9573c9a2760c964ed1659cc1673; path=/;
HttpOnly; Secure; SameSite=None
cache-control: private
```

11.3.12. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.

Procedure

- Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

- 1** For **<installation_directory>**, specify the location of your customized **./install-config.yaml** file.

- 2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to `<installation_directory>/openshift_install.log`.



IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

11.3.13. Configuring the default storage container

After you install the cluster, you must install the Nutanix CSI Operator and configure the default storage container for the cluster.

For more information, see the Nutanix documentation for [installing the CSI Operator](#) and [configuring registry storage](#).

11.3.14. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

11.3.15. Additional resources

- [About remote health monitoring](#)

11.3.16. Next steps

- [Opt out of remote health reporting](#)
- [Customize your cluster](#)

11.4. INSTALLING A CLUSTER ON NUTANIX IN A RESTRICTED NETWORK

In OpenShift Container Platform 4.16, you can install a cluster on Nutanix infrastructure in a restricted network by creating an internal mirror of the installation release content.

11.4.1. Prerequisites

- You have reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- The installation program requires access to port 9440 on Prism Central and Prism Element. You verified that port 9440 is accessible.
- If you use a firewall, you have met these prerequisites:
 - You confirmed that port 9440 is accessible. Control plane nodes must be able to reach Prism Central and Prism Element on port 9440 for the installation to succeed.
 - You configured the firewall to [grant access](#) to the sites that OpenShift Container Platform requires. This includes the use of Telemetry.
- If your Nutanix environment is using the default self-signed SSL/TLS certificate, replace it with a certificate that is signed by a CA. The installation program requires a valid CA-signed certificate to access to the Prism Central API. For more information about replacing the self-signed certificate, see the [Nutanix AOS Security Guide](#).


If your Nutanix environment uses an internal CA to issue certificates, you must configure a cluster-wide proxy as part of the installation process. For more information, see [Configuring a custom PKI](#).



IMPORTANT

Use 2048-bit certificates. The installation fails if you use 4096-bit certificates with Prism Central 2022.x.

- You have a container image registry, such as Red Hat Quay. If you do not already have a registry, you can create a mirror registry using [mirror registry for Red Hat OpenShift](#).
- You have used the [oc-mirror OpenShift CLI \(oc\) plugin](#) to mirror all of the required OpenShift Container Platform content and other images, including the Nutanix CSI Operator, to your mirror registry.



IMPORTANT

Because the installation media is on the mirror host, you can use that computer to complete all installation steps.

11.4.2. About installations in restricted networks

In OpenShift Container Platform 4.16, you can perform an installation that does not require an active connection to the internet to obtain software components. Restricted network installations can be completed using installer-provisioned infrastructure or user-provisioned infrastructure, depending on the cloud platform to which you are installing the cluster.

If you choose to perform a restricted network installation on a cloud platform, you still require access to its cloud APIs. Some cloud functions, like Amazon Web Service's Route 53 DNS and IAM services, require internet access. Depending on your network, you might require less internet access for an installation on bare metal hardware, Nutanix, or on VMware vSphere.

To complete a restricted network installation, you must create a registry that mirrors the contents of the OpenShift image registry and contains the installation media. You can create this registry on a mirror host, which can access both the internet and your closed network, or by using other methods that meet your restrictions.

11.4.2.1. Additional limits

Clusters in restricted networks have the following additional limitations and restrictions:


- The **ClusterVersion** status includes an **Unable to retrieve available updates** error.
- By default, you cannot use the contents of the Developer Catalog because you cannot access the required image stream tags.

11.4.3. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.

**IMPORTANT**

Do not skip this procedure in production environments, where disaster recovery and debugging is required.

**NOTE**

You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.

**NOTE**

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.

**NOTE**

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```

**NOTE**

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1 Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

11.4.4. Adding Nutanix root CA certificates to your system trust

Because the installation program requires access to the Prism Central API, you must add your Nutanix trusted root CA certificates to your system trust before you install an OpenShift Container Platform cluster.

Procedure

1. From the Prism Central web console, download the Nutanix root CA certificates.
2. Extract the compressed file that contains the Nutanix root CA certificates.
3. Add the files for your operating system to the system trust. For example, on a Fedora operating system, run the following command:

```
# cp certs/lin/* /etc/pki/ca-trust/source/anchors
```

4. Update your system trust. For example, on a Fedora operating system, run the following command:

```
# update-ca-trust extract
```

11.4.5. Downloading the RHCOS cluster image

Prism Central requires access to the Red Hat Enterprise Linux CoreOS (RHCOS) image to install the cluster. You can use the installation program to locate and download the RHCOS image and make it available through an internal HTTP server or Nutanix Objects.

Prerequisites

- Obtain the OpenShift Container Platform installation program and the pull secret for your cluster. For a restricted network installation, these files are on your mirror host.

Procedure

1. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install coreos print-stream-json
```

2. Use the output of the command to find the location of the Nutanix image, and click the link to download it.

Example output

```
"nutanix": {
  "release": "411.86.202210041459-0",
  "formats": {
    "qcow2": {
      "disk": {
        "location": "https://rhcos.mirror.openshift.com/art/storage/releases/rhcos-4.11/411.86.202210041459-0/x86_64/rhcos-411.86.202210041459-0-nutanix.x86_64.qcow2",
        "sha256":
"42e227cac6f11ac37ee8a2f9528bb3665146566890577fd55f9b950949e5a54b"
```

3. Make the image available through an internal HTTP server or Nutanix Objects.
4. Note the location of the downloaded image. You update the **platform** section in the installation configuration file (**install-config.yaml**) with the image's location before deploying the cluster.

Snippet of an `install-config.yaml` file that specifies the RHCOS image

```
platform:
  nutanix:
    clusterOSImage: http://example.com/images/rhcos-411.86.202210041459-0-nutanix.x86_64.qcow2
```

11.4.6. Creating the installation configuration file

You can customize the OpenShift Container Platform cluster you install on Nutanix.

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster. For a restricted network installation, these files are on your mirror host.
- You have the **imageContentSourcePolicy.yaml** file that was created when you mirrored your registry.
- You have the location of the Red Hat Enterprise Linux CoreOS (RHCOS) image you download.
- You have obtained the contents of the certificate for your mirror registry.

- You have retrieved a Red Hat Enterprise Linux CoreOS (RHCOS) image and uploaded it to an accessible location.
- You have verified that you have met the Nutanix networking requirements. For more information, see "Preparing to install on Nutanix".

Procedure

1. Create the **install-config.yaml** file.
 - a. Change to the directory that contains the installation program and run the following command:

```
┌ $ ./openshift-install create install-config --dir <installation_directory> 1
```

- 1 For **<installation_directory>**, specify the directory name to store the files that the installation program creates.

When specifying the directory:

- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
- Use an empty directory. Some installation assets, such as bootstrap X.509 certificates, have short expiration intervals, therefore you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

- b. At the prompts, provide the configuration details for your cloud:
 - i. Optional: Select an SSH key to use to access your cluster machines.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- ii. Select **nutanix** as the platform to target.
- iii. Enter the Prism Central domain name or IP address.
- iv. Enter the port that is used to log into Prism Central.
- v. Enter the credentials that are used to log into Prism Central.
The installation program connects to Prism Central.
- vi. Select the Prism Element that will manage the OpenShift Container Platform cluster.
- vii. Select the network subnet to use.
- viii. Enter the virtual IP address that you configured for control plane API access.

- ix. Enter the virtual IP address that you configured for cluster ingress.
 - x. Enter the base domain. This base domain must be the same one that you configured in the DNS records.
 - xi. Enter a descriptive name for your cluster.
The cluster name you enter must match the cluster name you specified when configuring the DNS records.
2. In the **install-config.yaml** file, set the value of **platform.nutanix.clusterOSImage** to the image location or name. For example:

```
platform:
  nutanix:
    clusterOSImage: http://mirror.example.com/images/rhcos-47.83.202103221318-0-
    nutanix.x86_64.qcow2
```

3. Edit the **install-config.yaml** file to give the additional information that is required for an installation in a restricted network.

- a. Update the **pullSecret** value to contain the authentication information for your registry:

```
pullSecret: '{"auths":{"<mirror_host_name>:5000": {"auth": "<credentials>","email":
"you@example.com"}}}'
```

For **<mirror_host_name>**, specify the registry domain name that you specified in the certificate for your mirror registry, and for **<credentials>**, specify the base64-encoded user name and password for your mirror registry.

- b. Add the **additionalTrustBundle** parameter and value.

```
additionalTrustBundle: |
  -----BEGIN CERTIFICATE-----

  //////////////////////////////////////
  -----END CERTIFICATE-----
```

The value must be the contents of the certificate file that you used for your mirror registry. The certificate file can be an existing, trusted certificate authority, or the self-signed certificate that you generated for the mirror registry.

- c. Add the image content resources, which resemble the following YAML excerpt:

```
imageContentSources:
- mirrors:
  - <mirror_host_name>:5000/<repo_name>/release
    source: quay.io/openshift-release-dev/ocp-release
- mirrors:
  - <mirror_host_name>:5000/<repo_name>/release
    source: registry.redhat.io/ocp/release
```

For these values, use the **imageContentSourcePolicy.yaml** file that was created when you mirrored the registry.

- d. Optional: Set the publishing strategy to **Internal**:

-

publish: Internal

By setting this option, you create an internal Ingress Controller and a private load balancer.


- Optional: Update one or more of the default configuration parameters in the **install.config.yaml** file to customize the installation.
For more information about the parameters, see "Installation configuration parameters".



NOTE

If you are installing a three-node cluster, be sure to set the **compute.replicas** parameter to **0**. This ensures that cluster's control planes are schedulable. For more information, see "Installing a three-node cluster on {platform}".

- Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT


The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

Additional resources

- [Installation configuration parameters for Nutanix](#)

11.4.6.1. Sample customized install-config.yaml file for Nutanix

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.



IMPORTANT

This sample YAML file is provided for reference only. You must obtain your **install-config.yaml** file by using the installation program and modify it.

```

apiVersion: v1
baseDomain: example.com 1
compute: 2
- hyperthreading: Enabled 3
  name: worker
  replicas: 3
platform:
  nutanix: 4
    cpus: 2
    coresPerSocket: 2
    memoryMiB: 8196
    osDisk:
      diskSizeGiB: 120
  categories: 5
    - key: <category_key_name>
      value: <category_value>
controlPlane: 6

```

```

hyperthreading: Enabled 7
name: master
replicas: 3
platform:
  nutanix: 8
    cpus: 4
    coresPerSocket: 2
    memoryMiB: 16384
    osDisk:
      diskSizeGiB: 120
    categories: 9
      - key: <category_key_name>
        value: <category_value>
metadata:
  creationTimestamp: null
  name: test-cluster 10
networking:
  clusterNetwork:
    - cidr: 10.128.0.0/14
      hostPrefix: 23
  machineNetwork:
    - cidr: 10.0.0.0/16
  networkType: OVNKubernetes 11
  serviceNetwork:
    - 172.30.0.0/16
platform:
  nutanix:
    apiVIP: 10.40.142.7 12
    ingressVIP: 10.40.142.8 13
  defaultMachinePlatform:
    bootType: Legacy
    categories: 14
      - key: <category_key_name>
        value: <category_value>
    project: 15
      type: name
      name: <project_name>
  prismCentral:
    endpoint:
      address: your.prismcentral.domainname 16
      port: 9440 17
    password: <password> 18
    username: <username> 19
  prismElements:
    - endpoint:
      address: your.prismelement.domainname
      port: 9440
      uuid: 0005b0f1-8f43-a0f2-02b7-3cecef193712
    subnetUUIDs:
      - c7938dc6-7659-453e-a688-e26020c68e43
    clusterOSImage: http://example.com/images/rhcos-47.83.202103221318-0-nutanix.x86_64.qcow2
20
credentialsMode: Manual
publish: External

```


`registry.example.com:5000`. For `<credentials>`, specify the base64-encoded user name and password for your mirror registry.

- 22 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.



IMPORTANT

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 23 Optional: You can provide the `sshKey` value that you use to access the machines in your cluster.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your `ssh-agent` process uses.

- 24 Provide the contents of the certificate file that you used for your mirror registry.
- 25 Provide these values from the `metadata.name: release-0` section of the `imageContentSourcePolicy.yaml` file that was created when you mirrored the registry.

11.4.6.2. Configuring failure domains

Failure domains improve the fault tolerance of an OpenShift Container Platform cluster by distributing control plane and compute machines across multiple Nutanix Prism Elements (clusters).

TIP

It is recommended that you configure three failure domains to ensure high-availability.

Prerequisites

- You have an installation configuration file (`install-config.yaml`).

Procedure

- Edit the `install-config.yaml` file and add the following stanza to configure the first failure domain:

```
apiVersion: v1
baseDomain: example.com
compute:
# ...
platform:
  nutanix:
```

```

failureDomains:
- name: <failure_domain_name>
  prismElement:
    name: <prism_element_name>
    uuid: <prism_element_uuid>
  subnetUUIDs:
  - <network_uuid>
# ...

```

where:

<failure_domain_name>

Specifies a unique name for the failure domain. The name is limited to 64 or fewer characters, which can include lower-case letters, digits, and a dash (-). The dash cannot be in the leading or ending position of the name.

<prism_element_name>

Optional. Specifies the name of the Prism Element.

<prism_element_uuid>

Specifies the UUID of the Prism Element.

<network_uuid>

Specifies the UUID of the Prism Element subnet object. The subnet's IP address prefix (CIDR) should contain the virtual IP addresses that the OpenShift Container Platform cluster uses. Only one subnet per failure domain (Prism Element) in an OpenShift Container Platform cluster is supported.

2. As required, configure additional failure domains.
3. To distribute control plane and compute machines across the failure domains, do one of the following:
 - If compute and control plane machines can share the same set of failure domains, add the failure domain names under the cluster's default machine configuration.

Example of control plane and compute machines sharing a set of failure domains

```

apiVersion: v1
baseDomain: example.com
compute:
# ...
platform:
  nutanix:
    defaultMachinePlatform:
      failureDomains:
      - failure-domain-1
      - failure-domain-2
      - failure-domain-3
# ...

```

- If compute and control plane machines must use different failure domains, add the failure domain names under the respective machine pools.

Example of control plane and compute machines using different failure domains


```

apiVersion: v1
baseDomain: example.com
compute:
# ...
controlPlane:
  platform:
    nutanix:
      failureDomains:
        - failure-domain-1
        - failure-domain-2
        - failure-domain-3
# ...
compute:
  platform:
    nutanix:
      failureDomains:
        - failure-domain-1
        - failure-domain-2
# ...

```

4. Save the file.

11.4.6.3. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```

apiVersion: v1
baseDomain: my.domain.com


```

```

proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5

```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.



NOTE

The installation program does not support the proxy **readinessEndpoints** field.



NOTE

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.


The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.

**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

11.4.7. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.

**IMPORTANT**

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**.
To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.

3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.



NOTE

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your PATH.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

11.4.8. Configuring IAM for Nutanix

Installing the cluster requires that the Cloud Credential Operator (CCO) operate in manual mode. While the installation program configures the CCO for manual mode, you must specify the identity and access management secrets.

Prerequisites

- You have configured the **ccoctl** binary.
- You have an **install-config.yaml** file.

Procedure

1. Create a YAML file that contains the credentials data in the following format:

Credentials data format

```
credentials:
- type: basic_auth 1
  data:
    prismCentral: 2
      username: <username_for_prism_central>
      password: <password_for_prism_central>
    prismElements: 3
      - name: <name_of_prism_element>
        username: <username_for_prism_element>
        password: <password_for_prism_element>
```

- 1 Specify the authentication type. Only basic authentication is supported.
 - 2 Specify the Prism Central credentials.
 - 3 Optional: Specify the Prism Element credentials.
2. Set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

3. Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included 1 \
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml 2 \
  --to=<path_to_directory_for_credentials_requests> 3
```

- 1 The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2 Specify the location of the **install-config.yaml** file.
- 3 Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

Sample CredentialsRequest object

```

apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  annotations:
    include.release.openshift.io/self-managed-high-availability: "true"
  labels:
    controller-tools.k8s.io: "1.0"
  name: openshift-machine-api-nutanix
  namespace: openshift-cloud-credential-operator
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: NutanixProviderSpec
  secretRef:
    name: nutanix-credentials
    namespace: openshift-machine-api

```

4. Use the **ccoctl** tool to process all **CredentialsRequest** objects by running the following command:

```

$ ccoctl nutanix create-shared-secrets \
  --credentials-requests-dir=<path_to_credentials_requests_directory> \ 1
  --output-dir=<ccoctl_output_dir> \ 2
  --credentials-source-filepath=<path_to_credentials_file> \ 3

```

- 1** Specify the path to the directory that contains the files for the component **CredentialsRequests** objects.
- 2** Optional: Specify the directory in which you want the **ccoctl** utility to create objects. By default, the utility creates objects in the directory in which the commands are run.
- 3** Optional: Specify the directory that contains the credentials data YAML file. By default, **ccoctl** expects this file to be in **<home_directory>/./nutanix/credentials**.

5. Edit the **install-config.yaml** configuration file so that the **credentialsMode** parameter is set to **Manual**.

Example **install-config.yaml** configuration file

```

apiVersion: v1
baseDomain: cluster1.example.com
credentialsMode: Manual 1
...

```

- 1** Add this line to set the **credentialsMode** parameter to **Manual**.

6. Create the installation manifests by running the following command:

```

$ openshift-install create manifests --dir <installation_directory> 1

```

- 1** Specify the path to the directory that contains the **install-config.yaml** file for your cluster.

- Copy the generated credential files to the target manifests directory by running the following command:

```
$ cp <ccoctl_output_dir>/manifests/*credentials.yaml ./<installation_directory>/manifests
```

Verification

- Ensure that the appropriate secrets exist in the **manifests** directory.

```
$ ls ./<installation_directory>/manifests
```

Example output

```
cluster-config.yaml
cluster-dns-02-config.yml
cluster-infrastructure-02-config.yml
cluster-ingress-02-config.yml
cluster-network-01-crd.yml
cluster-network-02-config.yml
cluster-proxy-01-config.yaml
cluster-scheduler-02-config.yml
cvo-overrides.yaml
kube-cloud-config.yaml
kube-system-configmap-root-ca.yaml
machine-config-server-tls-secret.yaml
openshift-config-secret-pull-secret.yaml
openshift-cloud-controller-manager-nutanix-credentials-credentials.yaml
openshift-machine-api-nutanix-credentials-credentials.yaml
```

11.4.9. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.

Procedure

- Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

- 1** For **<installation_directory>**, specify the location of your customized **./install-config.yaml** file.
- 2** To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to **<installation_directory>/openshift_install.log**.



IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

11.4.10. Post installation

Complete the following steps to complete the configuration of your cluster.

11.4.10.1. Disabling the default OperatorHub catalog sources

Operator catalogs that source content provided by Red Hat and community projects are configured for OperatorHub by default during an OpenShift Container Platform installation. In a restricted network environment, you must disable the default catalogs as a cluster administrator.

Procedure

- Disable the sources for the default catalogs by adding **disableAllDefaultSources: true** to the **OperatorHub** object:

```
$ oc patch OperatorHub cluster --type json \
  -p '[{"op": "add", "path": "/spec/disableAllDefaultSources", "value": true}]'
```

TIP

Alternatively, you can use the web console to manage catalog sources. From the **Administration** → **Cluster Settings** → **Configuration** → **OperatorHub** page, click the **Sources** tab, where you can create, update, delete, disable, and enable individual sources.

11.4.10.2. Installing the policy resources into the cluster

Mirroring the OpenShift Container Platform content using the oc-mirror OpenShift CLI (oc) plugin creates resources, which include **catalogSource-certified-operator-index.yaml** and **imageContentSourcePolicy.yaml**.

- The **ImageContentSourcePolicy** resource associates the mirror registry with the source registry and redirects image pull requests from the online registries to the mirror registry.
- The **CatalogSource** resource is used by Operator Lifecycle Manager (OLM) to retrieve information about the available Operators in the mirror registry, which lets users discover and install Operators.

After you install the cluster, you must install these resources into the cluster.

Prerequisites

- You have mirrored the image set to the registry mirror in the disconnected environment.
- You have access to the cluster as a user with the **cluster-admin** role.

Procedure

1. Log in to the OpenShift CLI as a user with the **cluster-admin** role.
2. Apply the YAML files from the results directory to the cluster:

```
$ oc apply -f ./oc-mirror-workspace/results-<id>/
```

Verification

1. Verify that the **ImageContentSourcePolicy** resources were successfully installed:

```
$ oc get imagecontentsourcepolicy
```

2. Verify that the **CatalogSource** resources were successfully installed:

```
$ oc get catalogsource --all-namespaces
```

11.4.10.3. Configuring the default storage container

After you install the cluster, you must install the Nutanix CSI Operator and configure the default storage container for the cluster.

For more information, see the Nutanix documentation for [installing the CSI Operator](#) and [configuring registry storage](#).

11.4.11. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

11.4.12. Additional resources

- [About remote health monitoring](#)

11.4.13. Next steps

- If necessary, see [Opt out of remote health reporting](#)
- If necessary, see [Registering your disconnected cluster](#)
- [Customize your cluster](#)

11.5. INSTALLING A THREE-NODE CLUSTER ON NUTANIX

In OpenShift Container Platform version 4.16, you can install a three-node cluster on Nutanix. A three-node cluster consists of three control plane machines, which also act as compute machines. This type of cluster provides a smaller, more resource efficient cluster, for cluster administrators and developers to use for testing, development, and production.

11.5.1. Configuring a three-node cluster

You configure a three-node cluster by setting the number of worker nodes to **0** in the **install-config.yaml** file before deploying the cluster. Setting the number of worker nodes to **0** ensures that the control plane machines are schedulable. This allows application workloads to be scheduled to run from the control plane nodes.



NOTE

Because application workloads run from control plane nodes, additional subscriptions are required, as the control plane nodes are considered to be compute nodes.

Prerequisites

- You have an existing **install-config.yaml** file.

Procedure

- Set the number of compute replicas to **0** in your **install-config.yaml** file, as shown in the following **compute** stanza:

Example **install-config.yaml** file for a three-node cluster

```
apiVersion: v1
baseDomain: example.com
compute:
- name: worker
  platform: {}
  replicas: 0
# ...
```

11.5.2. Next steps

- [Installing a cluster on Nutanix](#)

11.6. UNINSTALLING A CLUSTER ON NUTANIX

You can remove a cluster that you deployed to Nutanix.

11.6.1. Removing a cluster that uses installer-provisioned infrastructure

You can remove a cluster that uses installer-provisioned infrastructure from your cloud.



NOTE

After uninstallation, check your cloud provider for any resources not removed properly, especially with User Provisioned Infrastructure (UPI) clusters. There might be resources that the installer did not create or that the installer is unable to access.

Prerequisites

- You have a copy of the installation program that you used to deploy the cluster.
- You have the files that the installation program generated when you created your cluster.

Procedure

1. From the directory that contains the installation program on the computer that you used to install the cluster, run the following command:

```
$ ./openshift-install destroy cluster \
--dir <installation_directory> --log-level info 1 2
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

- 2 To view different details, specify **warn**, **debug**, or **error** instead of **info**.



NOTE

You must specify the directory that contains the cluster definition files for your cluster. The installation program requires the **metadata.json** file in this directory to delete the cluster.

- Optional: Delete the **<installation_directory>** directory and the OpenShift Container Platform installation program.

11.7. INSTALLATION CONFIGURATION PARAMETERS FOR NUTANIX

Before you deploy an OpenShift Container Platform cluster on Nutanix, you provide parameters to customize your cluster and the platform that hosts it. When you create the **install-config.yaml** file, you provide values for the required parameters through the command line. You can then modify the **install-config.yaml** file to customize your cluster further.

11.7.1. Available installation configuration parameters for Nutanix

The following tables specify the required, optional, and Nutanix-specific installation configuration parameters that you can set as part of the installation process.



NOTE

After installation, you cannot modify these parameters in the **install-config.yaml** file.

11.7.1.1. Required configuration parameters

Required installation configuration parameters are described in the following table:

Table 11.3. Required parameters

Parameter	Description	Values
apiVersion:	The API version for the install-config.yaml content. The current version is v1 . The installation program may also support older API versions.	String

Parameter	Description	Values
baseDomain:	The base domain of your cloud provider. The base domain is used to create routes to your OpenShift Container Platform cluster components. The full DNS name for your cluster is a combination of the baseDomain and metadata.name parameter values that uses the <metadata.name>.<baseDomain> format.	A fully-qualified domain or subdomain name, such as example.com .
metadata:	Kubernetes resource ObjectMeta , from which only the name parameter is consumed.	Object
metadata: name:	The name of the cluster. DNS records for the cluster are all subdomains of {{.metadata.name}} . {{.baseDomain}} .	String of lowercase letters and hyphens (-), such as dev .
platform:	The configuration for the specific platform upon which to perform the installation: aws, baremetal, azure, gcp, ibmcloud, nutanix, openstack, powersvs, vsphere , or {} . For additional information about platform . <platform> parameters, consult the table for your specific platform that follows.	Object

Parameter	Description	Values
<code>pullSecret:</code>	Get a pull secret from Red Hat OpenShift Cluster Manager to authenticate downloading container images for OpenShift Container Platform components from services such as Quay.io.	<pre>{ "auths":{ "cloud.openshift.com":{ "auth":"b3Blb=", "email":"you@example.com" }, "quay.io":{ "auth":"b3Blb=", "email":"you@example.com" } } }</pre>

11.7.1.2. Network configuration parameters

You can customize your installation configuration based on the requirements of your existing network infrastructure. For example, you can expand the IP address block for the cluster network or provide different IP address blocks than the defaults.


Only IPv4 addresses are supported.




NOTE

Globalnet is not supported with Red Hat OpenShift Data Foundation disaster recovery solutions. For regional disaster recovery scenarios, ensure that you use a nonoverlapping range of private IP addresses for the cluster and service networks in each cluster.

Table 11.4. Network parameters

Parameter	Description	Values
<code>networking:</code>	The configuration for the cluster network.	Object  NOTE You cannot modify parameters specified by the networking object after installation.
<code>networking: networkType:</code>	The Red Hat OpenShift Networking network plugin to install.	OVNKubernetes. OVNKubernetes is a CNI plugin for Linux networks and hybrid networks that contain both Linux and Windows servers. The default value is OVNKubernetes .


Parameter	Description	Values
<code>networking: clusterNetwork:</code>	<p>The IP address blocks for pods.</p> <p>The default value is 10.128.0.0/14 with a host prefix of /23.</p> <p>If you specify multiple IP address blocks, the blocks must not overlap.</p>	<p>An array of objects. For example:</p> <pre>networking: clusterNetwork: - cidr: 10.128.0.0/14 hostPrefix: 23</pre>
<code>networking: clusterNetwork: cidr:</code>	<p>Required if you use networking.clusterNetwork. An IP address block.</p> <p>An IPv4 network.</p>	<p>An IP address block in Classless Inter-Domain Routing (CIDR) notation. The prefix length for an IPv4 block is between 0 and 32.</p>
<code>networking: clusterNetwork: hostPrefix:</code>	<p>The subnet prefix length to assign to each individual node. For example, if hostPrefix is set to 23 then each node is assigned a /23 subnet out of the given cidr. A hostPrefix value of 23 provides 510 ($2^{(32 - 23)} - 2$) pod IP addresses.</p>	<p>A subnet prefix.</p> <p>The default value is 23.</p>
<code>networking: serviceNetwork:</code>	<p>The IP address block for services. The default value is 172.30.0.0/16.</p> <p>The OVN-Kubernetes network plugins supports only a single IP address block for the service network.</p>	<p>An array with an IP address block in CIDR format. For example:</p> <pre>networking: serviceNetwork: - 172.30.0.0/16</pre>
<code>networking: machineNetwork:</code>	<p>The IP address blocks for machines.</p> <p>If you specify multiple IP address blocks, the blocks must not overlap.</p>	<p>An array of objects. For example:</p> <pre>networking: machineNetwork: - cidr: 10.0.0.0/16</pre>
<code>networking: machineNetwork: cidr:</code>	<p>Required if you use networking.machineNetwork. An IP address block. The default value is 10.0.0.0/16 for all platforms other than libvirt and IBM Power® Virtual Server. For libvirt, the default value is 192.168.126.0/24. For IBM Power® Virtual Server, the default value is 192.168.0.0/24.</p>	<p>An IP network block in CIDR notation.</p> <p>For example, 10.0.0.0/16.</p> <div style="display: flex; align-items: flex-start;">  <div> <p>NOTE</p> <p>Set the networking.machineNetwork to match the CIDR that the preferred NIC resides in.</p> </div> </div>

11.7.1.3. Optional configuration parameters

Optional installation configuration parameters are described in the following table:

Table 11.5. Optional parameters



Parameter	Description	Values
<code>additionalTrustBundle:</code>	A PEM-encoded X.509 certificate bundle that is added to the nodes' trusted certificate store. This trust bundle may also be used when a proxy has been configured.	String
<code>capabilities:</code>	Controls the installation of optional core cluster components. You can reduce the footprint of your OpenShift Container Platform cluster by disabling optional components. For more information, see the "Cluster capabilities" page in <i>Installing</i> .	String array
<code>capabilities: baselineCapabilitySet:</code>	Selects an initial set of optional capabilities to enable. Valid values are None , v4.11 , v4.12 and vCurrent . The default value is vCurrent .	String
<code>capabilities: additionalEnabledCapabilities:</code>	Extends the set of optional capabilities beyond what you specify in baselineCapabilitySet . You may specify multiple capabilities in this parameter.	String array
<code>cpuPartitioningMode:</code>	Enables workload partitioning, which isolates OpenShift Container Platform services, cluster management workloads, and infrastructure pods to run on a reserved set of CPUs. Workload partitioning can only be enabled during installation and cannot be disabled after installation. While this field enables workload partitioning, it does not configure workloads to use specific CPUs. For more information, see the <i>Workload partitioning</i> page in the <i>Scalability and Performance</i> section.	None or AllNodes . None is the default value.
<code>compute:</code>	The configuration for the machines that comprise the compute nodes.	Array of MachinePool objects.

Parameter	Description	Values
<code>compute: architecture:</code>	Determines the instruction set architecture of the machines in the pool. Currently, clusters with varied architectures are not supported. All pools must specify the same architecture. Valid values are amd64 (the default).	String
<code>compute: hyperthreading:</code>	<p>Whether to enable or disable simultaneous multithreading, or hyperthreading, on compute machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.</p> <div style="display: flex; align-items: center;">  <div> <p>IMPORTANT</p> <p>If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.</p> </div> </div>	Enabled or Disabled
<code>compute: name:</code>	Required if you use compute . The name of the machine pool.	worker
<code>compute: platform:</code>	Required if you use compute . Use this parameter to specify the cloud provider to host the worker machines. This parameter value must match the controlPlane.platform parameter value.	aws, azure, gcp, ibmcloud, nutanix, openstack, powervs, vsphere , or {}
<code>compute: replicas:</code>	The number of compute machines, which are also known as worker machines, to provision.	A positive integer greater than or equal to 2 . The default value is 3 .

Parameter	Description	Values
<code>featureSet:</code>	Enables the cluster for a feature set. A feature set is a collection of OpenShift Container Platform features that are not enabled by default. For more information about enabling a feature set during installation, see "Enabling features using feature gates".	String. The name of the feature set to enable, such as TechPreviewNoUpgrade .
<code>controlPlane:</code>	The configuration for the machines that comprise the control plane.	Array of MachinePool objects.
<code>controlPlane: architecture:</code>	Determines the instruction set architecture of the machines in the pool. Currently, clusters with varied architectures are not supported. All pools must specify the same architecture. Valid values are amd64 (the default).	String
<code>controlPlane: hyperthreading:</code>	<p>Whether to enable or disable simultaneous multithreading, or hyperthreading, on control plane machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.</p> <div style="display: flex; align-items: flex-start;"> <div style="background-color: black; width: 20px; height: 100px; margin-right: 10px;"></div> <div> <p>IMPORTANT</p> <p>If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.</p> </div> </div>	Enabled or Disabled
<code>controlPlane: name:</code>	Required if you use controlPlane . The name of the machine pool.	master
<code>controlPlane: platform:</code>	Required if you use controlPlane . Use this parameter to specify the cloud provider that hosts the control plane machines. This parameter value must match the compute.platform parameter value.	aws, azure, gcp, ibmcloud, nutanix, openstack, powervs, vsphere , or {}

Parameter	Description	Values
controlPlane: replicas:	The number of control plane machines to provision.	Supported values are 3 , or 1 when deploying single-node OpenShift.
credentialsMode:	The Cloud Credential Operator (CCO) mode. If no mode is specified, the CCO dynamically tries to determine the capabilities of the provided credentials, with a preference for mint mode on the platforms where multiple modes are supported.	Mint, Passthrough, Manual or an empty string (""). ^[1]

Parameter	Description	Values
fips:	<p>Enable or disable FIPS mode. The default is false (disabled). If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.</p> <div style="display: flex; align-items: flex-start;"> <div style="background-color: black; width: 20px; height: 100%; margin-right: 10px;"></div> <div> <p>IMPORTANT</p> <p>To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see Installing the system in FIPS mode.</p> <p>When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.</p> </div> </div> <div style="display: flex; align-items: flex-start; margin-top: 10px;"> <div style="background-color: black; width: 20px; height: 100%; margin-right: 10px;"></div> <div> <p>NOTE</p> <p>If you are using Azure File storage, you cannot enable FIPS mode.</p> </div> </div>	false or true
imageContentSources:	Sources and repositories for the release-image content.	Array of objects. Includes a source and, optionally, mirrors , as described in the following rows of this table.

Parameter	Description	Values
<code>imageContentSources: source:</code>	Required if you use imageContentSources . Specify the repository that users refer to, for example, in image pull specifications.	String
<code>imageContentSources: mirrors:</code>	Specify one or more repositories that may also contain the same images.	Array of strings
<code>publish:</code>	How to publish or expose the user-facing endpoints of your cluster, such as the Kubernetes API, OpenShift routes.	<p>Internal or External. The default value is External.</p> <p>Setting this field to Internal is not supported on non-cloud platforms.</p> <div style="display: flex; align-items: flex-start;">  <div> <p>IMPORTANT</p> <p>If the value of the field is set to Internal, the cluster will become non-functional. For more information, refer to BZ#1953035.</p> </div> </div>
<code>sshKey:</code>	<p>The SSH key to authenticate access to your cluster machines.</p> <div style="display: flex; align-items: flex-start;">  <div> <p>NOTE</p> <p>For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your ssh-agent process uses.</p> </div> </div>	For example, sshKey: ssh-ed25519 AAAA...

1. Not all CCO modes are supported for all cloud providers. For more information about CCO modes, see the "Managing cloud provider credentials" entry in the *Authentication and authorization* content.

11.7.1.4. Additional Nutanix configuration parameters

Additional Nutanix configuration parameters are described in the following table:

Table 11.6. Additional Nutanix cluster parameters

Parameter	Description	Values
compute: platform: nutanix: categories: key:	The name of a prism category key to apply to compute VMs. This parameter must be accompanied by the value parameter, and both key and value parameters must exist in Prism Central. For more information on categories, see Category management .	String
compute: platform: nutanix: categories: value:	The value of a prism category key-value pair to apply to compute VMs. This parameter must be accompanied by the key parameter, and both key and value parameters must exist in Prism Central.	String
compute: platform: nutanix: failureDomains:	The failure domains that apply to only compute machines. Failure domains are specified in platform.nutanix.failureDomains .	List. The name of one or more failures domains.
compute: platform: nutanix: project: type:	The type of identifier you use to select a project for compute VMs. Projects define logical groups of user roles for managing permissions, networks, and other parameters. For more information on projects, see Projects Overview .	name or uuid
compute: platform: nutanix: project: name: or uuid:	The name or UUID of a project with which compute VMs are associated. This parameter must be accompanied by the type parameter.	String
compute: platform: nutanix: bootType:	The boot type that the compute machines use. You must use the Legacy boot type in OpenShift Container Platform 4.16. For more information on boot types, see Understanding UEFI, Secure Boot, and TPM in the Virtualized Environment .	Legacy , SecureBoot or UEFI . The default is Legacy .

Parameter	Description	Values
controlPlane: platform: nutanix: categories: key:	The name of a prism category key to apply to control plane VMs. This parameter must be accompanied by the value parameter, and both key and value parameters must exist in Prism Central. For more information on categories, see Category management .	String
controlPlane: platform: nutanix: categories: value:	The value of a prism category key-value pair to apply to control plane VMs. This parameter must be accompanied by the key parameter, and both key and value parameters must exist in Prism Central.	String
controlPlane: platform: nutanix: failureDomains:	The failure domains that apply to only control plane machines. Failure domains are specified in platform.nutanix.failureDomains .	List. The name of one or more failures domains.
controlPlane: platform: nutanix: project: type:	The type of identifier you use to select a project for control plane VMs. Projects define logical groups of user roles for managing permissions, networks, and other parameters. For more information on projects, see Projects Overview .	name or uuid
controlPlane: platform: nutanix: project: name: or uuid:	The name or UUID of a project with which control plane VMs are associated. This parameter must be accompanied by the type parameter.	String
platform: nutanix: defaultMachinePlatform: categories: key:	The name of a prism category key to apply to all VMs. This parameter must be accompanied by the value parameter, and both key and value parameters must exist in Prism Central. For more information on categories, see Category management .	String

Parameter	Description	Values
<pre>platform: nutanix: defaultMachinePlatform: categories: value:</pre>	<p>The value of a prism category key-value pair to apply to all VMs. This parameter must be accompanied by the key parameter, and both key and value parameters must exist in Prism Central.</p>	String
<pre>platform: nutanix: defaultMachinePlatform: failureDomains:</pre>	<p>The failure domains that apply to both control plane and compute machines.</p> <p>Failure domains are specified in platform.nutanix.failureDomains.</p>	<p>List.</p> <p>The name of one or more failures domains.</p>
<pre>platform: nutanix: defaultMachinePlatform: project: type:</pre>	<p>The type of identifier you use to select a project for all VMs. Projects define logical groups of user roles for managing permissions, networks, and other parameters. For more information on projects, see Projects Overview.</p>	name or uuid .
<pre>platform: nutanix: defaultMachinePlatform: project: name: or uuid:</pre>	<p>The name or UUID of a project with which all VMs are associated. This parameter must be accompanied by the type parameter.</p>	String
<pre>platform: nutanix: defaultMachinePlatform: bootType:</pre>	<p>The boot type for all machines. You must use the Legacy boot type in OpenShift Container Platform 4.16. For more information on boot types, see Understanding UEFI, Secure Boot, and TPM in the Virtualized Environment.</p>	Legacy , SecureBoot or UEFI . The default is Legacy .

Parameter	Description	Values
platform: nutanix: apiVIP:	The virtual IP (VIP) address that you configured for control plane API access.	IP address
platform: nutanix: failureDomains: - name: prismElement: name: uuid: subnetUUIDs: -	By default, the installation program installs cluster machines to a single Prism Element instance. You can specify additional Prism Element instances for fault tolerance, and then apply them to: <ul style="list-style-type: none"> • The cluster's default machine configuration • Only control plane or compute machine pools 	A list of configured failure domains. For more information on usage, see "Configuring a failure domain" in "Installing a cluster on Nutanix".
platform: nutanix: ingressVIP:	The virtual IP (VIP) address that you configured for cluster ingress.	IP address
platform: nutanix: prismCentral: endpoint: address:	The Prism Central domain name or IP address.	String
platform: nutanix: prismCentral: endpoint: port:	The port that is used to log into Prism Central.	String
platform: nutanix: prismCentral: password:	The password for the Prism Central user name.	String

Parameter	Description	Values
platform: nutanix: prismCentral: username:	The user name that is used to log into Prism Central.	String
platform: nutanix: prismElements: endpoint: address:	The Prism Element domain name or IP address. [1]	String
platform: nutanix: prismElements: endpoint: port:	The port that is used to log into Prism Element.	String
platform: nutanix: prismElements: uuid:	The universally unique identifier (UUID) for Prism Element.	String
platform: nutanix: subnetUUIDs:	The UUID of the Prism Element network that contains the virtual IP addresses and DNS records that you configured. [2]	String
platform: nutanix: clusterOSImage:	Optional: By default, the installation program downloads and installs the Red Hat Enterprise Linux CoreOS (RHCOS) image. If Prism Central does not have internet access, you can override the default behavior by hosting the RHCOS image on any HTTP server and pointing the installation program to the image.	An HTTP or HTTPS URL, optionally with a SHA-256 checksum. For example, <code>http://example.com/images/rhcos-47.83.202103221318-0-nutanix.x86_64.qcow2</code>

1. The **prismElements** section holds a list of Prism Elements (clusters). A Prism Element encompasses all of the Nutanix resources, for example virtual machines and subnets, that are used to host the OpenShift Container Platform cluster.
2. Only one subnet per Prism Element in an OpenShift Container Platform cluster is supported.

CHAPTER 12. INSTALLING ON-PREMISE WITH ASSISTED INSTALLER

12.1. INSTALLING AN ON-PREMISE CLUSTER USING THE ASSISTED INSTALLER

You can install OpenShift Container Platform on on-premise hardware or on-premise VMs by using the Assisted Installer. Installing OpenShift Container Platform by using the Assisted Installer supports **x86_64**, **AArch64**, **ppc64le**, and **s390x** CPU architectures.

12.1.1. Using the Assisted Installer

The [Assisted Installer](#) is a user-friendly installation solution offered on the [Red Hat Hybrid Cloud Console](#). The Assisted Installer supports the various deployment platforms with a focus on bare metal, Nutanix, and vSphere infrastructures.


The Assisted Installer provides installation functionality as a service. This software-as-a-service (SaaS) approach has the following advantages:

- **Web user interface:** The web user interface performs cluster installation without the user having to create the installation configuration files manually.
- **No bootstrap node:** A bootstrap node is not required when installing with the Assisted Installer. The bootstrapping process executes on a node within the cluster.
- **Hosting:** The Assisted Installer hosts:
 - Ignition files
 - The installation configuration
 - A discovery ISO
 - The installer
- **Streamlined installation workflow:** Deployment does not require in-depth knowledge of OpenShift Container Platform. The Assisted Installer provides reasonable defaults and provides the installer as a service, which:
 - Eliminates the need to install and run the OpenShift Container Platform installer locally.
 - Ensures the latest version of the installer up to the latest tested z-stream releases. Older versions remain available, if needed.
 - Enables building automation by using the API without the need to run the OpenShift Container Platform installer locally.
- **Advanced networking:** The Assisted Installer supports IPv4 networking with SDN and OVN, IPv6 and dual stack networking with OVN only, NMState-based static IP addressing, and an HTTP/S proxy. OVN is the default Container Network Interface (CNI) for OpenShift Container Platform 4.12 and later releases. SDN is supported up to OpenShift Container Platform 4.14, but is not supported for OpenShift Container Platform 4.15 and later releases.
- **Preinstallation validation:** The Assisted Installer validates the configuration before installation to ensure a high probability of success. The validation process includes the following checks:

- Ensuring network connectivity
 - Ensuring sufficient network bandwidth
 - Ensuring connectivity to the registry
 - Ensuring time synchronization between cluster nodes
 - Verifying that the cluster nodes meet the minimum hardware requirements
 - Validating the installation configuration parameters
- **REST API:** The Assisted Installer has a REST API, enabling automation.

The Assisted Installer supports installing OpenShift Container Platform on premises in a connected environment, including with an optional HTTP/S proxy. It can install the following:

- Highly available OpenShift Container Platform or single-node OpenShift (SNO)
- OpenShift Container Platform on bare metal, Nutanix, or vSphere with full platform integration, or other virtualization platforms without integration
- Optional: OpenShift Virtualization, multicluster engine, Logical Volume Manager (LVM) Storage, and OpenShift Data Foundation

**NOTE**

Currently, OpenShift Virtualization and LVM Storage are not supported on IBM Z® (**s390x**) architecture.

The user interface provides an intuitive interactive workflow where automation does not exist or is not required. Users may also automate installations using the REST API.

See the [Assisted Installer for OpenShift Container Platform](#) documentation for details.

12.1.2. API support for the Assisted Installer

Supported APIs for the Assisted Installer are stable for a minimum of three months from the announcement of deprecation.

CHAPTER 13. INSTALLING AN ON-PREMISE CLUSTER WITH THE AGENT-BASED INSTALLER

13.1. PREPARING TO INSTALL WITH THE AGENT-BASED INSTALLER

13.1.1. About the Agent-based Installer

The Agent-based installation method provides the flexibility to boot your on-premises servers in any way that you choose. It combines the ease of use of the Assisted Installation service with the ability to run offline, including in air-gapped environments. Agent-based installation is a subcommand of the OpenShift Container Platform installer. It generates a bootable ISO image containing all of the information required to deploy an OpenShift Container Platform cluster, with an available release image.

The configuration is in the same format as for the installer-provisioned infrastructure and user-provisioned infrastructure installation methods. The Agent-based Installer can also optionally generate or accept Zero Touch Provisioning (ZTP) custom resources. ZTP allows you to provision new edge sites with declarative configurations of bare-metal equipment.

Table 13.1. Agent-based Installer supported architectures

CPU architecture	Connected installation	Disconnected installation
64-bit x86	✓	✓
64-bit ARM	✓	✓
ppc64le	✓	✓
s390x	✓	✓

13.1.2. Understanding Agent-based Installer

As an OpenShift Container Platform user, you can leverage the advantages of the Assisted Installer hosted service in disconnected environments.

The Agent-based installation comprises a bootable ISO that contains the Assisted discovery agent and the Assisted Service. Both are required to perform the cluster installation, but the latter runs on only one of the hosts.

The **openshift-install agent create image** subcommand generates an ephemeral ISO based on the inputs that you provide. You can choose to provide inputs through the following manifests:

Preferred:

- **install-config.yaml**
- **agent-config.yaml**

Optional: ZTP manifests

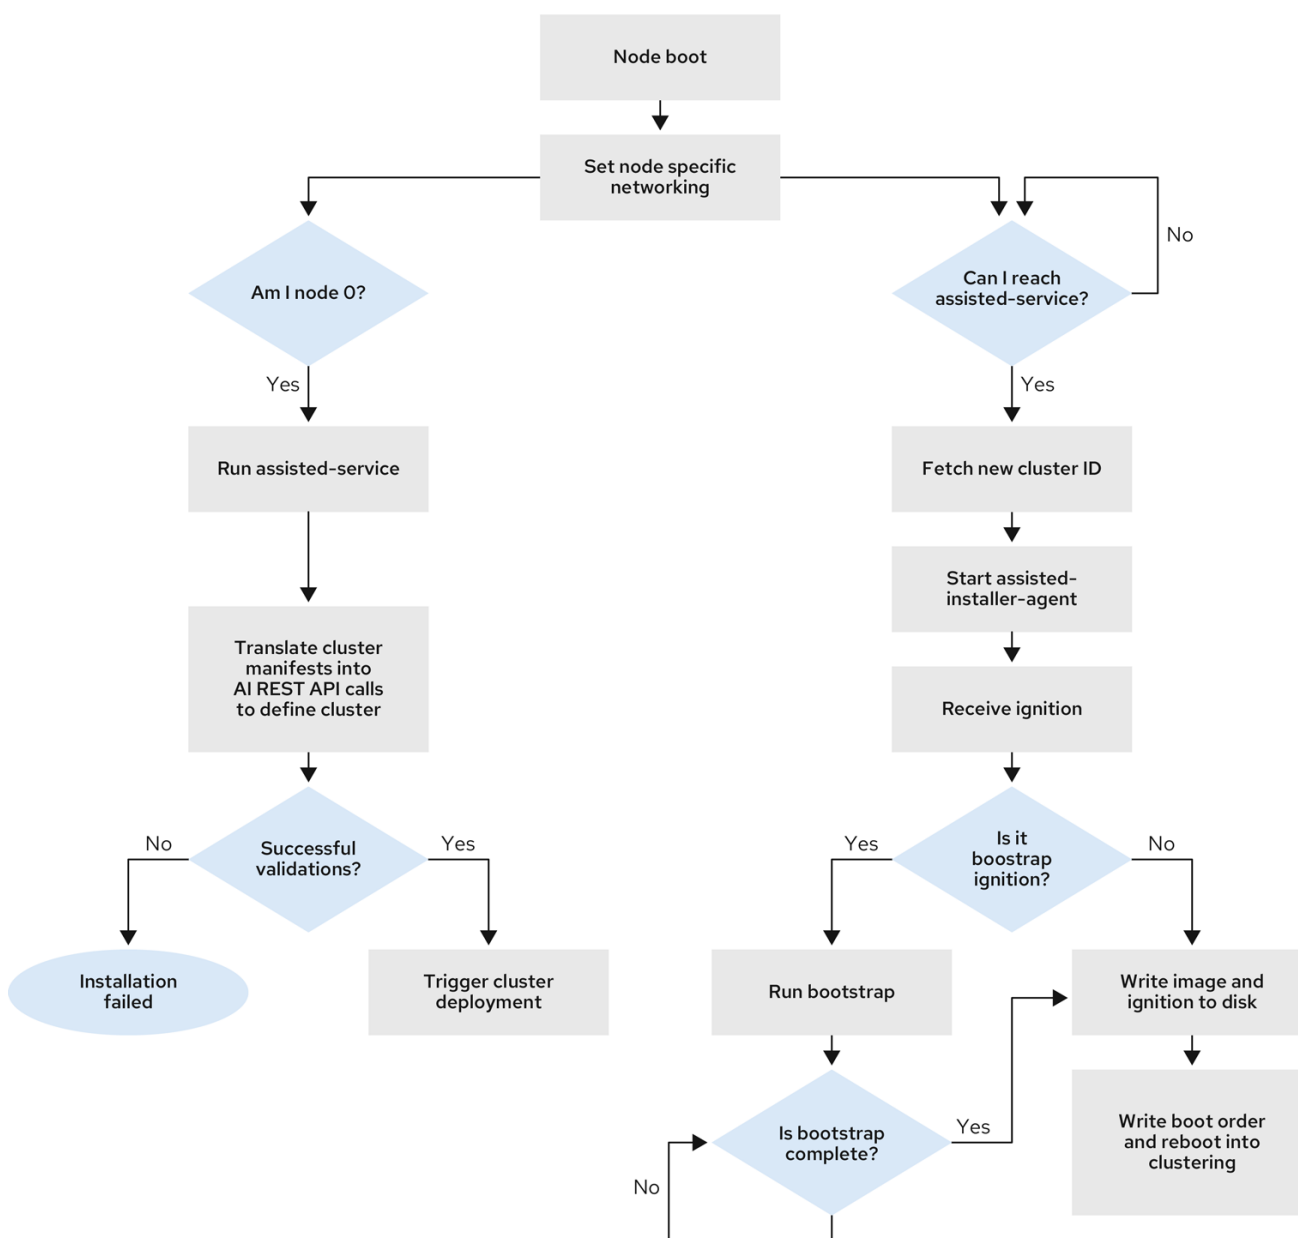
- **cluster-manifests/cluster-deployment.yaml**

- **cluster-manifests/agent-cluster-install.yaml**
- **cluster-manifests/pull-secret.yaml**
- **cluster-manifests/infraenv.yaml**
- **cluster-manifests/cluster-image-set.yaml**
- **cluster-manifests/nmstateconfig.yaml**
- **mirror/registries.conf**
- **mirror/ca-bundle.crt**

13.1.2.1. Agent-based Installer workflow

One of the control plane hosts runs the Assisted Service at the start of the boot process and eventually becomes the bootstrap host. This node is called the **rendezvous host** (node 0). The Assisted Service ensures that all the hosts meet the requirements and triggers an OpenShift Container Platform cluster deployment. All the nodes have the Red Hat Enterprise Linux CoreOS (RHCOS) image written to the disk. The non-bootstrap nodes reboot and initiate a cluster deployment. Once the nodes are rebooted, the rendezvous host reboots and joins the cluster. The bootstrapping is complete and the cluster is deployed.

Figure 13.1. Node installation workflow



281_OpenShift_1022

You can install a disconnected OpenShift Container Platform cluster through the **openshift-install agent create image** subcommand for the following topologies:

- **A single-node OpenShift Container Platform cluster (SNO)** A node that is both a master and worker.
- **A three-node OpenShift Container Platform cluster:** A compact cluster that has three master nodes that are also worker nodes.
- **Highly available OpenShift Container Platform cluster (HA)** Three master nodes with any number of worker nodes.

13.1.2.2. Recommended resources for topologies

Recommended cluster resources for the following topologies:

Table 13.2. Recommended cluster resources

Topology	Number of control plane nodes	Number of compute nodes	vCPU	Memory	Storage
Single-node cluster	1	0	8 vCPU cores	16 GB of RAM	120 GB
Compact cluster	3	0 or 1	8 vCPU cores	16 GB of RAM	120 GB
HA cluster	3	2 and above	8 vCPU cores	16 GB of RAM	120 GB

In the `install-config.yaml`, specify the platform on which to perform the installation. The following platforms are supported:

- **baremetal**
- **vsphere**
- **none**



IMPORTANT

For platform **none**:

- The **none** option requires the provision of DNS name resolution and load balancing infrastructure in your cluster. See *Requirements for a cluster using the platform "none" option* in the "Additional resources" section for more information.
- Review the information in the [guidelines for deploying OpenShift Container Platform on non-tested platforms](#) before you attempt to install an OpenShift Container Platform cluster in virtualized or cloud environments.

Additional resources

- [Requirements for a cluster using the platform "none" option](#)
- [Increase the network MTU](#)
- [Adding worker nodes to single-node OpenShift clusters](#)

13.1.3. About FIPS compliance

For many OpenShift Container Platform customers, regulatory readiness, or compliance, on some level is required before any systems can be put into production. That regulatory readiness can be imposed by national standards, industry standards or the organization's corporate governance framework. Federal Information Processing Standards (FIPS) compliance is one of the most critical components required in highly secure environments to ensure that only supported cryptographic technologies are allowed on nodes.



IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

13.1.4. Configuring FIPS through the Agent-based Installer

During a cluster deployment, the Federal Information Processing Standards (FIPS) change is applied when the Red Hat Enterprise Linux CoreOS (RHCOS) machines are deployed in your cluster. For Red Hat Enterprise Linux (RHEL) machines, you must enable FIPS mode when you install the operating system on the machines that you plan to use as worker machines.

You can enable FIPS mode through the preferred method of **install-config.yaml** and **agent-config.yaml**:

1. You must set value of the **fips** field to **True** in the **install-config.yaml** file:

Sample install-config.yaml file

```
apiVersion: v1
baseDomain: test.example.com
metadata:
  name: sno-cluster
fips: True
```

2. Optional: If you are using the GitOps ZTP manifests, you must set the value of **fips** as **True** in the **Agent-install.openshift.io/install-config-overrides** field in the **agent-cluster-install.yaml** file:

Sample agent-cluster-install.yaml file


```
apiVersion: extensions.hive.openshift.io/v1beta1
kind: AgentClusterInstall
metadata:
  annotations:
    agent-install.openshift.io/install-config-overrides: '{"fips": True}'
  name: sno-cluster
  namespace: sno-cluster-test
```

Additional resources

- [OpenShift Security Guide Book](#)
- [Support for FIPS cryptography](#)

13.1.5. Host configuration

You can make additional configurations for each host on the cluster in the **agent-config.yaml** file, such as network configurations and root device hints.



IMPORTANT

For each host you configure, you must provide the MAC address of an interface on the host to specify which host you are configuring.

13.1.5.1. Host roles

Each host in the cluster is assigned a role of either **master** or **worker**. You can define the role for each host in the **agent-config.yaml** file by using the **role** parameter. If you do not assign a role to the hosts, the roles will be assigned at random during installation.

It is recommended to explicitly define roles for your hosts.

The **rendezvousIP** must be assigned to a host with the **master** role. This can be done manually or by allowing the Agent-based Installer to assign the role.



IMPORTANT

You do not need to explicitly define the **master** role for the rendezvous host, however you cannot create configurations that conflict with this assignment.

For example, if you have 4 hosts with 3 of the hosts explicitly defined to have the **master** role, the last host that is automatically assigned the **worker** role during installation cannot be configured as the rendezvous host.

Sample agent-config.yaml file

```
apiVersion: v1beta1
kind: AgentConfig
metadata:
  name: example-cluster
rendezvousIP: 192.168.111.80
hosts:
- hostname: master-1
  role: master
  interfaces:
  - name: eno1
    macAddress: 00:ef:44:21:e6:a5
- hostname: master-2
  role: master
  interfaces:
  - name: eno1
    macAddress: 00:ef:44:21:e6:a6
- hostname: master-3
  role: master
  interfaces:
  - name: eno1
    macAddress: 00:ef:44:21:e6:a7
- hostname: worker-1
  role: worker
```

```

interfaces:
- name: eno1
  macAddress: 00:ef:44:21:e6:a8

```

13.1.5.2. About root device hints

The **rootDeviceHints** parameter enables the installer to provision the Red Hat Enterprise Linux CoreOS (RHCOS) image to a particular device. The installer examines the devices in the order it discovers them, and compares the discovered values with the hint values. The installer uses the first discovered device that matches the hint value. The configuration can combine multiple hints, but a device must match all hints for the installer to select it.

Table 13.3. Subfields

Subfield	Description
deviceName	A string containing a Linux device name such as /dev/vda or /dev/disk/by-path/ . It is recommended to use the /dev/disk/by-path/<device_path> link to the storage location. The hint must match the actual value exactly.
hctl	A string containing a SCSI bus address like 0:0:0:0 . The hint must match the actual value exactly.
model	A string containing a vendor-specific device identifier. The hint can be a substring of the actual value.
vendor	A string containing the name of the vendor or manufacturer of the device. The hint can be a substring of the actual value.
serialNumber	A string containing the device serial number. The hint must match the actual value exactly.
minSizeGigabytes	An integer representing the minimum size of the device in gigabytes.
wwn	A string containing the unique storage identifier. The hint must match the actual value exactly.
rotational	A boolean indicating whether the device should be a rotating disk (true) or not (false).

Example usage

```

- name: master-0
  role: master
  rootDeviceHints:
    deviceName: "/dev/sda"

```

13.1.6. About networking

The **rendezvous IP** must be known at the time of generating the agent ISO, so that during the initial boot all the hosts can check in to the assisted service. If the IP addresses are assigned using a Dynamic Host Configuration Protocol (DHCP) server, then the **rendezvousIP** field must be set to an IP address of one of the hosts that will become part of the deployed control plane. In an environment without a DHCP server, you can define IP addresses statically.

In addition to static IP addresses, you can apply any network configuration that is in NMState format. This includes VLANs and NIC bonds.

13.1.6.1. DHCP

Preferred method: `install-config.yaml` and `agent-config.yaml`

You must specify the value for the **rendezvousIP** field. The **networkConfig** fields can be left blank:

Sample `agent-config.yaml` file

```
apiVersion: v1alpha1
kind: AgentConfig
metadata:
  name: sno-cluster
rendezvousIP: 192.168.111.80 1
```

- 1 The IP address for the rendezvous host.

13.1.6.2. Static networking

- a. Preferred method: `install-config.yaml` and `agent-config.yaml`

Sample `agent-config.yaml` file

```
cat > agent-config.yaml << EOF
apiVersion: v1alpha1
kind: AgentConfig
metadata:
  name: sno-cluster
rendezvousIP: 192.168.111.80 1
hosts:
- hostname: master-0
  interfaces:
  - name: eno1
    macAddress: 00:ef:44:21:e6:a5 2
networkConfig:
  interfaces:
  - name: eno1
    type: ethernet
    state: up
    mac-address: 00:ef:44:21:e6:a5
    ipv4:
      enabled: true
      address:
      - ip: 192.168.111.80 3
```

```

    prefix-length: 23 4
    dhcp: false
  dns-resolver:
    config:
      server:
        - 192.168.111.1 5
  routes:
    config:
      - destination: 0.0.0.0/0
        next-hop-address: 192.168.111.1 6
        next-hop-interface: eno1
        table-id: 254
EOF

```

- 1 If a value is not specified for the **rendezvousIP** field, one address will be chosen from the static IP addresses specified in the **networkConfig** fields.
- 2 The MAC address of an interface on the host, used to determine which host to apply the configuration to.
- 3 The static IP address of the target bare metal host.
- 4 The static IP address's subnet prefix for the target bare metal host.
- 5 The DNS server for the target bare metal host.
- 6 Next hop address for the node traffic. This must be in the same subnet as the IP address set for the specified interface.

b. Optional method: GitOps ZTP manifests

The optional method of the GitOps ZTP custom resources comprises 6 custom resources; you can configure static IPs in the **nmstateconfig.yaml** file.

```

apiVersion: agent-install.openshift.io/v1beta1
kind: NMStateConfig
metadata:
  name: master-0
  namespace: openshift-machine-api
  labels:
    cluster0-nmstate-label-name: cluster0-nmstate-label-value
spec:
  config:
    interfaces:
      - name: eth0
        type: ethernet
        state: up
        mac-address: 52:54:01:aa:aa:a1
        ipv4:
          enabled: true
          address:
            - ip: 192.168.122.2 1
              prefix-length: 23 2
          dhcp: false
    dns-resolver:

```

```

config:
  server:
    - 192.168.122.1 3
  routes:
    config:
      - destination: 0.0.0.0/0
        next-hop-address: 192.168.122.1 4
        next-hop-interface: eth0
        table-id: 254
    interfaces:
      - name: eth0
        macAddress: 52:54:01:aa:aa:a1 5

```

- 1** The static IP address of the target bare metal host.
- 2** The static IP address's subnet prefix for the target bare metal host.
- 3** The DNS server for the target bare metal host.
- 4** Next hop address for the node traffic. This must be in the same subnet as the IP address set for the specified interface.
- 5** The MAC address of an interface on the host, used to determine which host to apply the configuration to.

The rendezvous IP is chosen from the static IP addresses specified in the **config** fields.

13.1.7. Requirements for a cluster using the platform "none" option

This section describes the requirements for an Agent-based OpenShift Container Platform installation that is configured to use the platform **none** option.



IMPORTANT

Review the information in the [guidelines for deploying OpenShift Container Platform on non-tested platforms](#) before you attempt to install an OpenShift Container Platform cluster in virtualized or cloud environments.

13.1.7.1. Platform "none" DNS requirements

In OpenShift Container Platform deployments, DNS name resolution is required for the following components:

- The Kubernetes API
- The OpenShift Container Platform application wildcard
- The control plane and compute machines

Reverse DNS resolution is also required for the Kubernetes API, the control plane machines, and the compute machines.

DNS A/AAAA or CNAME records are used for name resolution and PTR records are used for reverse name resolution. The reverse records are important because Red Hat Enterprise Linux CoreOS

(RHCOS) uses the reverse records to set the hostnames for all the nodes, unless the hostnames are provided by DHCP. Additionally, the reverse records are used to generate the certificate signing requests (CSR) that OpenShift Container Platform needs to operate.




NOTE

It is recommended to use a DHCP server to provide the hostnames to each cluster node.

The following DNS records are required for an OpenShift Container Platform cluster using the platform **none** option and they must be in place before installation. In each record, **<cluster_name>** is the cluster name and **<base_domain>** is the base domain that you specify in the **install-config.yaml** file. A complete DNS record takes the form: **<component>.<cluster_name>.<base_domain>..**

Table 13.4. Required DNS records

Component	Record	Description
Kubernetes API	api.<cluster_name>.<base_domain>.	A DNS A/AAAA or CNAME record, and a DNS PTR record, to identify the API load balancer. These records must be resolvable by both clients external to the cluster and from all the nodes within the cluster.
	api-int.<cluster_name>.<base_domain>.	A DNS A/AAAA or CNAME record, and a DNS PTR record, to internally identify the API load balancer. These records must be resolvable from all the nodes within the cluster.
		 <h3>IMPORTANT</h3> <p>The API server must be able to resolve the worker nodes by the hostnames that are recorded in Kubernetes. If the API server cannot resolve the node names, then proxied API calls can fail, and you cannot retrieve logs from pods.</p>
Routes	*.apps.<cluster_name>.<base_domain>.	A wildcard DNS A/AAAA or CNAME record that refers to the application ingress load balancer. The application ingress load balancer targets the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default. These records must be resolvable by both clients external to the cluster and from all the nodes within the cluster. For example, console-openshift-console.apps.<cluster_name>.<base_domain> is used as a wildcard route to the OpenShift Container Platform console.
Control plane machines	<master><n>.<cluster_name>.<base_domain>.	DNS A/AAAA or CNAME records and DNS PTR records to identify each machine for the control plane nodes. These records must be resolvable by the nodes within the cluster.

Component	Record	Description
Computes machine s	<worker><n>. <cluster_name>. <base_domain>.	DNS A/AAAA or CNAME records and DNS PTR records to identify each machine for the worker nodes. These records must be resolvable by the nodes within the cluster.

**NOTE**

In OpenShift Container Platform 4.4 and later, you do not need to specify etcd host and SRV records in your DNS configuration.

TIP

You can use the **dig** command to verify name and reverse name resolution.

13.1.7.1.1. Example DNS configuration for platform "none" clusters

This section provides A and PTR record configuration samples that meet the DNS requirements for deploying OpenShift Container Platform using the platform **none** option. The samples are not meant to provide advice for choosing one DNS solution over another.

In the examples, the cluster name is **ocp4** and the base domain is **example.com**.

Example DNS A record configuration for a platform "none" cluster

The following example is a BIND zone file that shows sample A records for name resolution in a cluster using the platform **none** option.

Example 13.1. Sample DNS zone database

```
$TTL 1W
@ IN SOA ns1.example.com. root (
  2019070700 ; serial
  3H ; refresh (3 hours)
  30M ; retry (30 minutes)
  2W ; expiry (2 weeks)
  1W ) ; minimum (1 week)
IN NS ns1.example.com.
IN MX 10 smtp.example.com.
;
;
ns1.example.com. IN A 192.168.1.5
smtp.example.com. IN A 192.168.1.5
;
helper.example.com. IN A 192.168.1.5
helper.ocp4.example.com. IN A 192.168.1.5
;
api.ocp4.example.com. IN A 192.168.1.5 1
api-int.ocp4.example.com. IN A 192.168.1.5 2
;
```



```
*.apps.ocp4.example.com. IN A 192.168.1.5 3
;
master0.ocp4.example.com. IN A 192.168.1.97 4
master1.ocp4.example.com. IN A 192.168.1.98 5
master2.ocp4.example.com. IN A 192.168.1.99 6
;
worker0.ocp4.example.com. IN A 192.168.1.11 7
worker1.ocp4.example.com. IN A 192.168.1.7 8
;
;EOF
```

- 1 Provides name resolution for the Kubernetes API. The record refers to the IP address of the API load balancer.
- 2 Provides name resolution for the Kubernetes API. The record refers to the IP address of the API load balancer and is used for internal cluster communications.
- 3 Provides name resolution for the wildcard routes. The record refers to the IP address of the application ingress load balancer. The application ingress load balancer targets the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.



NOTE

In the example, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

- 4 5 6 Provides name resolution for the control plane machines.
- 7 8 Provides name resolution for the compute machines.

Example DNS PTR record configuration for a platform "none" cluster

The following example BIND zone file shows sample PTR records for reverse name resolution in a cluster using the platform **none** option.

Example 13.2. Sample DNS zone database for reverse records

```
$TTL 1W
@ IN SOA ns1.example.com. root (
  2019070700 ; serial
  3H ; refresh (3 hours)
  30M ; retry (30 minutes)
  2W ; expiry (2 weeks)
  1W ) ; minimum (1 week)
IN NS ns1.example.com.
;
5.1.168.192.in-addr.arpa. IN PTR api.ocp4.example.com. 1
5.1.168.192.in-addr.arpa. IN PTR api-int.ocp4.example.com. 2
;
```

```

97.1.168.192.in-addr.arpa. IN PTR master0.ocp4.example.com. 3
98.1.168.192.in-addr.arpa. IN PTR master1.ocp4.example.com. 4
99.1.168.192.in-addr.arpa. IN PTR master2.ocp4.example.com. 5
;
11.1.168.192.in-addr.arpa. IN PTR worker0.ocp4.example.com. 6
7.1.168.192.in-addr.arpa. IN PTR worker1.ocp4.example.com. 7
;
;EOF

```

- 1 Provides reverse DNS resolution for the Kubernetes API. The PTR record refers to the record name of the API load balancer.
- 2 Provides reverse DNS resolution for the Kubernetes API. The PTR record refers to the record name of the API load balancer and is used for internal cluster communications.
- 3 4 5 Provides reverse DNS resolution for the control plane machines.
- 6 7 Provides reverse DNS resolution for the compute machines.

**NOTE**

A PTR record is not required for the OpenShift Container Platform application wildcard.

13.1.7.2. Platform "none" Load balancing requirements

Before you install OpenShift Container Platform, you must provision the API and application Ingress load balancing infrastructure. In production scenarios, you can deploy the API and application Ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

**NOTE**


These requirements do not apply to single-node OpenShift clusters using the platform **none** option.

**NOTE**

If you want to deploy the API and application Ingress load balancers with a Red Hat Enterprise Linux (RHEL) instance, you must purchase the RHEL subscription separately.

The load balancing infrastructure must meet the following requirements:

1. **API load balancer.** Provides a common endpoint for users, both human and machine, to interact with and configure the platform. Configure the following conditions:
 - Layer 4 load balancing only. This can be referred to as Raw TCP, SSL Passthrough, or SSL Bridge mode. If you use SSL Bridge mode, you must enable Server Name Indication (SNI) for the API routes.
 - A stateless load balancing algorithm. The options vary based on the load balancer implementation.

**IMPORTANT**

Do not configure session persistence for an API load balancer.

Configure the following ports on both the front and back of the load balancers:

Table 13.5. API load balancer

Port	Back-end machines (pool members)	Internal	External	Description
6443	Control plane. You must configure the /readyz endpoint for the API server health check probe.	X	X	Kubernetes API server
22623	Control plane.	X		Machine config server

**NOTE**

The load balancer must be configured to take a maximum of 30 seconds from the time the API server turns off the **/readyz** endpoint to the removal of the API server instance from the pool. Within the time frame after **/readyz** returns an error or becomes healthy, the endpoint must have been removed or added. Probing every 5 or 10 seconds, with two successful requests to become healthy and three to become unhealthy, are well-tested values.

2. **Application Ingress load balancer.** Provides an ingress point for application traffic flowing in from outside the cluster. A working configuration for the Ingress router is required for an OpenShift Container Platform cluster.

Configure the following conditions:

- Layer 4 load balancing only. This can be referred to as Raw TCP, SSL Passthrough, or SSL Bridge mode. If you use SSL Bridge mode, you must enable Server Name Indication (SNI) for the ingress routes.
- A connection-based or session-based persistence is recommended, based on the options available and types of applications that will be hosted on the platform.

TIP


If the true IP address of the client can be seen by the application Ingress load balancer, enabling source IP-based session persistence can improve performance for applications that use end-to-end TLS encryption.

Configure the following ports on both the front and back of the load balancers:

Table 13.6. Application Ingress load balancer

Port	Back-end machines (pool members)	Internal	External	Description
------	----------------------------------	----------	----------	-------------

Port	Back-end machines (pool members)	Internal	External	Description
443	The machines that run the Ingress Controller pods, compute, or worker, by default.	X	X	HTTPS traffic
80	The machines that run the Ingress Controller pods, compute, or worker, by default.	X	X	HTTP traffic

**NOTE**

If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application Ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes.

13.1.7.2.1. Example load balancer configuration for platform "none" clusters

This section provides an example API and application Ingress load balancer configuration that meets the load balancing requirements for clusters using the platform **none** option. The sample is an `/etc/haproxy/haproxy.cfg` configuration for an HAProxy load balancer. The example is not meant to provide advice for choosing one load balancing solution over another.

In the example, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

**NOTE**

If you are using HAProxy as a load balancer and SELinux is set to **enforcing**, you must ensure that the HAProxy service can bind to the configured TCP port by running **setsebool -P haproxy_connect_any=1**.

Example 13.3. Sample API and application Ingress load balancer configuration

```
global
  log      127.0.0.1 local2
  pidfile  /var/run/haproxy.pid
  maxconn  4000
  daemon
defaults
  mode          http
  log           global
  option        dontlognull
  option http-server-close
  option        redispatch
  retries       3
  timeout http-request  10s
  timeout queue    1m
```

```

timeout connect      10s
timeout client      1m
timeout server      1m
timeout http-keep-alive 10s
timeout check       10s
maxconn             3000
listen api-server-6443 1
bind *:6443
mode tcp
server master0 master0.ocp4.example.com:6443 check inter 1s
server master1 master1.ocp4.example.com:6443 check inter 1s
server master2 master2.ocp4.example.com:6443 check inter 1s
listen machine-config-server-22623 2
bind *:22623
mode tcp
server master0 master0.ocp4.example.com:22623 check inter 1s
server master1 master1.ocp4.example.com:22623 check inter 1s
server master2 master2.ocp4.example.com:22623 check inter 1s
listen ingress-router-443 3
bind *:443
mode tcp
balance source
server worker0 worker0.ocp4.example.com:443 check inter 1s
server worker1 worker1.ocp4.example.com:443 check inter 1s
listen ingress-router-80 4
bind *:80
mode tcp
balance source
server worker0 worker0.ocp4.example.com:80 check inter 1s
server worker1 worker1.ocp4.example.com:80 check inter 1s

```

- 1** Port **6443** handles the Kubernetes API traffic and points to the control plane machines.
- 2** Port **22623** handles the machine config server traffic and points to the control plane machines.
- 3** Port **443** handles the HTTPS traffic and points to the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.
- 4** Port **80** handles the HTTP traffic and points to the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.



NOTE

If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application Ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes.

TIP

If you are using HAProxy as a load balancer, you can check that the **haproxy** process is listening on ports **6443**, **22623**, **443**, and **80** by running **netstat -nltupe** on the HAProxy node.

13.1.8. Example: Bonds and VLAN interface node network configuration

The following **agent-config.yaml** file is an example of a manifest for bond and VLAN interfaces.

```
apiVersion: v1alpha1
kind: AgentConfig
rendezvousIP: 10.10.10.14
hosts:
- hostname: master0
  role: master
interfaces:
- name: enp0s4
  macAddress: 00:21:50:90:c0:10
- name: enp0s5
  macAddress: 00:21:50:90:c0:20
networkConfig:
  interfaces:
  - name: bond0.300 1
    type: vlan 2
    state: up
    vlan:
      base-iface: bond0
      id: 300
    ipv4:
      enabled: true
      address:
        - ip: 10.10.10.14
          prefix-length: 24
      dhcp: false
  - name: bond0 3
    type: bond 4
    state: up
    mac-address: 00:21:50:90:c0:10 5
    ipv4:
      enabled: false
    ipv6:
      enabled: false
    link-aggregation:
      mode: active-backup 6
      options:
        miimon: "150" 7
      port:
        - enp0s4
        - enp0s5
  dns-resolver: 8
    config:
      server:
        - 10.10.10.11
        - 10.10.10.12
  routes:
    config:
      - destination: 0.0.0.0/0
```

```

next-hop-address: 10.10.10.10 9
next-hop-interface: bond0.300 10
table-id: 254

```

- 1** **3** Name of the interface.
- 2** The type of interface. This example creates a VLAN.
- 4** The type of interface. This example creates a bond.
- 5** The mac address of the interface.
- 6** The **mode** attribute specifies the bonding mode.
- 7** Specifies the MII link monitoring frequency in milliseconds. This example inspects the bond link every 150 milliseconds.
- 8** Optional: Specifies the search and server settings for the DNS server.
- 9** Next hop address for the node traffic. This must be in the same subnet as the IP address set for the specified interface.
- 10** Next hop interface for the node traffic.

13.1.9. Example: Bonds and SR-IOV dual-nic node network configuration



IMPORTANT

Support for Day 1 operations associated with enabling NIC partitioning for SR-IOV devices is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

The following **agent-config.yaml** file is an example of a manifest for dual port NIC with a bond and SR-IOV interfaces:

```

apiVersion: v1alpha1
kind: AgentConfig
rendezvousIP: 10.10.10.14
hosts:
- hostname: worker-1
  interfaces:
  - name: eno1
    macAddress: 0c:42:a1:55:f3:06
  - name: eno2
    macAddress: 0c:42:a1:55:f3:07
networkConfig: 1
  interfaces: 2

```

```
- name: eno1 3
  type: ethernet 4
  state: up
  mac-address: 0c:42:a1:55:f3:06
  ipv4:
    enabled: true
    dhcp: false 5
  ethernet:
    sr-iov:
      total-vfs: 2 6
  ipv6:
    enabled: false
- name: sriov:eno1:0
  type: ethernet
  state: up 7
  ipv4:
    enabled: false 8
  ipv6:
    enabled: false
    dhcp: false
- name: sriov:eno1:1
  type: ethernet
  state: down
- name: eno2
  type: ethernet
  state: up
  mac-address: 0c:42:a1:55:f3:07
  ipv4:
    enabled: true
  ethernet:
    sr-iov:
      total-vfs: 2
  ipv6:
    enabled: false
- name: sriov:eno2:0
  type: ethernet
  state: up
  ipv4:
    enabled: false
  ipv6:
    enabled: false
- name: sriov:eno2:1
  type: ethernet
  state: down
- name: bond0
  type: bond
  state: up
  min-tx-rate: 100 9
  max-tx-rate: 200 10
  link-aggregation:
    mode: active-backup 11
    options:
      primary: sriov:eno1:0 12
  port:
    - sriov:eno1:0
```



```

- sriov:eno2:0
ipv4:
  address:
    - ip: 10.19.16.57 13
      prefix-length: 23
  dhcp: false
  enabled: true
ipv6:
  enabled: false
dns-resolver:
  config:
    server:
      - 10.11.5.160
      - 10.2.70.215
routes:
  config:
    - destination: 0.0.0.0/0
      next-hop-address: 10.19.17.254
      next-hop-interface: bond0 14
      table-id: 254

```

- 1** The **networkConfig** field contains information about the network configuration of the host, with subfields including **interfaces**, **dns-resolver**, and **routes**.
- 2** The **interfaces** field is an array of network interfaces defined for the host.
- 3** The name of the interface.
- 4** The type of interface. This example creates an ethernet interface.
- 5** Set this to **false** to disable DHCP for the physical function (PF) if it is not strictly required.
- 6** Set this to the number of SR-IOV virtual functions (VFs) to instantiate.
- 7** Set this to **up**.
- 8** Set this to **false** to disable IPv4 addressing for the VF attached to the bond.
- 9** Sets a minimum transmission rate, in Mbps, for the VF. This sample value sets a rate of 100 Mbps.
 - This value must be less than or equal to the maximum transmission rate.
 - Intel NICs do not support the **min-tx-rate** parameter. For more information, see [BZ#1772847](#).
- 10** Sets a maximum transmission rate, in Mbps, for the VF. This sample value sets a rate of 200 Mbps.
- 11** Sets the desired bond mode.
- 12** Sets the preferred port of the bonding interface. The primary device is the first of the bonding interfaces to be used and is not abandoned unless it fails. This setting is particularly useful when one NIC in the bonding interface is faster and, therefore, able to handle a bigger load. This setting is only valid when the bonding interface is in **active-backup** mode (mode 1) and **balance-tlb** (mode 5).
- 13** Sets a static IP address for the bond interface. This is the node IP address.

- 14 Sets **bond0** as the gateway for the default route.

Additional resources

- [Configuring network bonding](#)

13.1.10. Sample install-config.yaml file for bare metal

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.

```

apiVersion: v1
baseDomain: example.com 1
compute: 2
- name: worker
  replicas: 0 3
controlPlane: 4
  name: master
  replicas: 1 5
metadata:
  name: sno-cluster 6
networking:
  clusterNetwork:
  - cidr: 10.128.0.0/14 7
    hostPrefix: 23 8
  networkType: OVNKubernetes 9
  serviceNetwork: 10
  - 172.30.0.0/16
platform:
  none: {} 11
fips: false 12
pullSecret: '{"auths": ...}' 13
sshKey: 'ssh-ed25519 AAAA...' 14

```

- 1 The base domain of the cluster. All DNS records must be sub-domains of this base and include the cluster name.
- 2 4 The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.
- 3 This parameter controls the number of compute machines that the Agent-based installation waits to discover before triggering the installation process. It is the number of compute machines that must be booted with the generated ISO.



NOTE

If you are installing a three-node cluster, do not deploy any compute machines when you install the Red Hat Enterprise Linux CoreOS (RHCOS) machines.


- 5 The number of control plane machines that you add to the cluster. Because the cluster uses these values as the number of etcd endpoints in the cluster, the value must match the number of control
- 6 The cluster name that you specified in your DNS records.
- 7 A block of IP addresses from which pod IP addresses are allocated. This block must not overlap with existing physical networks. These IP addresses are used for the pod network. If you need to access the pods from an external network, you must configure load balancers and routers to manage the traffic.



NOTE

Class E CIDR range is reserved for a future use. To use the Class E CIDR range, you must ensure your networking environment accepts the IP addresses within the Class E CIDR range.

- 8 The subnet prefix length to assign to each individual node. For example, if **hostPrefix** is set to **23**, then each node is assigned a **/23** subnet out of the given **cidr**, which allows for 510 ($2^{(32 - 23)} - 2$) pod IP addresses. If you are required to provide access to nodes from an external network, configure load balancers and routers to manage the traffic.
- 9 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 10 The IP address pool to use for service IP addresses. You can enter only one IP address pool. This block must not overlap with existing physical networks. If you need to access the services from an external network, configure load balancers and routers to manage the traffic.
- 11 You must set the platform to **none** for a single-node cluster. You can set the platform to **vsphere**, **baremetal**, or **none** for multi-node clusters.



NOTE

If you set the platform to **vsphere** or **baremetal**, you can configure IP address endpoints for cluster nodes in three ways:

- IPv4
- IPv6
- IPv4 and IPv6 in parallel (dual-stack)

Example of dual-stack networking

```
networking:
  clusterNetwork:
    - cidr: 172.21.0.0/16
      hostPrefix: 23
    - cidr: fd02::/48
      hostPrefix: 64
  machineNetwork:
    - cidr: 192.168.11.0/16
    - cidr: 2001:DB8::/32
  serviceNetwork:
    - 172.22.0.0/16
    - fd03::/112
  networkType: OVNKubernetes
platform:
  baremetal:
    apiVIPs:
      - 192.168.11.3
      - 2001:DB8::4
    ingressVIPs:
      - 192.168.11.4
      - 2001:DB8::5
```


- 12** Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.



IMPORTANT

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 13** This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.
- 14** The SSH public key for the **core** user in Red Hat Enterprise Linux CoreOS (RHCOS).

**NOTE**

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

13.1.11. Validation checks before agent ISO creation

The Agent-based Installer performs validation checks on user defined YAML files before the ISO is created. Once the validations are successful, the agent ISO is created.

install-config.yaml

- **baremetal**, **vsphere** and **none** platforms are supported.
- The **networkType** parameter must be **OVNKubernetes** in the case of **none** platform.
- **apiVIPs** and **ingressVIPs** parameters must be set for bare metal and vSphere platforms.
- Some host-specific fields in the bare metal platform configuration that have equivalents in **agent-config.yaml** file are ignored. A warning message is logged if these fields are set.

agent-config.yaml

- Each interface must have a defined MAC address. Additionally, all interfaces must have a different MAC address.
- At least one interface must be defined for each host.
- World Wide Name (WWN) vendor extensions are not supported in root device hints.
- The **role** parameter in the **host** object must have a value of either **master** or **worker**.

13.1.11.1. ZTP manifests**agent-cluster-install.yaml**

- For IPv6, the only supported value for the **networkType** parameter is **OVNKubernetes**. The **OpenshiftSDN** value can be used only for IPv4.

cluster-image-set.yaml

- The **ReleaseImage** parameter must match the release defined in the installer.

13.1.12. Next steps

- [Installing a cluster with the Agent-based Installer](#)

13.2. UNDERSTANDING DISCONNECTED INSTALLATION MIRRORING

You can use a mirror registry for disconnected installations and to ensure that your clusters only use container images that satisfy your organization's controls on external content. Before you install a cluster on infrastructure that you provision in a disconnected environment, you must mirror the required

container images into that environment. To mirror container images, you must have a registry for mirroring.

13.2.1. Mirroring images for a disconnected installation through the Agent-based Installer

You can use one of the following procedures to mirror your OpenShift Container Platform image repository to your mirror registry:

- [Mirroring images for a disconnected installation](#)
- [Mirroring images for a disconnected installation using the oc-mirror plugin](#)

13.2.2. About mirroring the OpenShift Container Platform image repository for a disconnected registry

To use mirror images for a disconnected installation with the Agent-based Installer, you must modify the **install-config.yaml** file.

You can mirror the release image by using the output of either the **oc adm release mirror** or **oc mirror** command. This is dependent on which command you used to set up the mirror registry.

The following example shows the output of the **oc adm release mirror** command.

```
$ oc adm release mirror
```

Example output

To use the new mirrored repository to install, add the following section to the **install-config.yaml**:

```
imageContentSources:
  mirrors:
  virthost.ostest.test.metalkube.org:5000/localimages/local-release-image
  source: quay.io/openshift-release-dev/ocp-v4.0-art-dev
  mirrors:
  virthost.ostest.test.metalkube.org:5000/localimages/local-release-image
  source: registry.ci.openshift.org/ocp/release
```

The following example shows part of the **imageContentSourcePolicy.yaml** file generated by the **oc-mirror** plugin. The file can be found in the results directory, for example **oc-mirror-workspace/results-1682697932/**.

Example imageContentSourcePolicy.yaml file

```
spec:
  repositoryDigestMirrors:
  - mirrors:
    - virthost.ostest.test.metalkube.org:5000/openshift/release
    source: quay.io/openshift-release-dev/ocp-v4.0-art-dev
  - mirrors:
    - virthost.ostest.test.metalkube.org:5000/openshift/release-images
    source: quay.io/openshift-release-dev/ocp-release
```

13.2.2.1. Configuring the Agent-based Installer to use mirrored images

You must use the output of either the **oc adm release mirror** command or the **oc-mirror** plugin to configure the Agent-based Installer to use mirrored images.

Procedure

1. If you used the **oc-mirror** plugin to mirror your release images:
 - a. Open the **imageContentSourcePolicy.yaml** located in the results directory, for example **oc-mirror-workspace/results-1682697932/**.
 - b. Copy the text in the **repositoryDigestMirrors** section of the yaml file.
2. If you used the **oc adm release mirror** command to mirror your release images:
 - Copy the text in the **imageContentSources** section of the command output.
3. Paste the copied text into the **imageContentSources** field of the **install-config.yaml** file.
4. Add the certificate file used for the mirror registry to the **additionalTrustBundle** field of the yaml file.



IMPORTANT

The value must be the contents of the certificate file that you used for your mirror registry. The certificate file can be an existing, trusted certificate authority, or the self-signed certificate that you generated for the mirror registry.

Example install-config.yaml file

```
additionalTrustBundle: |
  -----BEGIN CERTIFICATE-----
  /-----END CERTIFICATE-----
```

5. If you are using GitOps ZTP manifests: add the **registries.conf** and **ca-bundle.crt** files to the **mirror** path to add the mirror configuration in the agent ISO image.



NOTE

You can create the **registries.conf** file from the output of either the **oc adm release mirror** command or the **oc mirror** plugin. The format of the **/etc/containers/registries.conf** file has changed. It is now version 2 and in TOML format.

Example registries.conf file

```
[[registry]]
location = "registry.ci.openshift.org/ocp/release" mirror-by-digest-only = true

[[registry.mirror]] location = "virthost.ostest.test.metalkube.org:5000/localimages/local-release-image"
```

```
[[registry]]  
location = "quay.io/openshift-release-dev/ocp-v4.0-art-dev" mirror-by-digest-only = true  
  
[[registry.mirror]] location = "virthost.ostest.test.metalkube.org:5000/localimages/local-  
release-image"
```

13.3. INSTALLING AN OPENSIFT CONTAINER PLATFORM CLUSTER WITH THE AGENT-BASED INSTALLER

Use the following procedures to install an OpenShift Container Platform cluster using the Agent-based Installer.

13.3.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- If you use a firewall or proxy, you [configured it to allow the sites](#) that your cluster requires access to.

13.3.2. Installing OpenShift Container Platform with the Agent-based Installer

The following procedures deploy a single-node OpenShift Container Platform in a disconnected environment. You can use these procedures as a basis and modify according to your requirements.

13.3.2.1. Downloading the Agent-based Installer

Use this procedure to download the Agent-based Installer and the CLI needed for your installation.

Procedure

1. Log in to the OpenShift Container Platform web console using your login credentials.
2. Navigate to [Datacenter](#).
3. Click **Run Agent-based Installer locally**.
4. Select the operating system and architecture for the **OpenShift Installer** and **Command line interface**.
5. Click **Download Installer** to download and extract the install program.
6. You can either download or copy the pull secret by clicking on **Download pull secret** or **Copy pull secret**.
7. Click **Download command-line tools** and place the **openshift-install** binary in a directory that is on your **PATH**.

13.3.2.2. Verifying the supported architecture for installing an Agent-based Installer cluster

Before installing an OpenShift Container Platform cluster using the Agent-based Installer, you can verify the supported architecture on which you can install the cluster. This procedure is optional.

Prerequisites

- You installed the OpenShift CLI (**oc**).
- You have downloaded the installation program.

Procedure

1. Log in to the OpenShift CLI (**oc**).
2. Check your release payload by running the following command:

```
$ ./openshift-install version
```

Example output

```
./openshift-install 4.16.0
built from commit abc123def456
release image quay.io/openshift-release-dev/ocp-
release@sha256:123abc456def789ghi012jkl345mno678pqr901stu234vwx567yz0
release architecture amd64
```

If you are using the release image with the **multi** payload, the **release architecture** displayed in the output of this command is the default architecture.

3. To check the architecture of the payload, run the following command:

```
$ oc adm release info <release_image> -o jsonpath="{ .metadata.metadata}" 1
```

- 1** Replace **<release_image>** with the release image. For example: **quay.io/openshift-release-dev/ocp-release@sha256:123abc456def789ghi012jkl345mno678pqr901stu234vwx567yz0**.

.Example output when the release image uses the **multi** payload:

```
{"release.openshift.io architecture":"multi"}
```

If you are using the release image with the **multi** payload, you can install the cluster on different architectures such as **arm64**, **amd64**, **s390x**, and **ppc64le**. Otherwise, you can install the cluster only on the **release architecture** displayed in the output of the **openshift-install version** command.

13.3.2.3. Creating the preferred configuration inputs

Use this procedure to create the preferred configuration inputs used to create the agent image.

Procedure

1. Install **nmstate** dependency by running the following command:

```
$ sudo dnf install /usr/bin/nmstatectl -y
```

- Place the **openshift-install** binary in a directory that is on your PATH.
- Create a directory to store the install configuration by running the following command:

```
$ mkdir ~/<directory_name>
```



NOTE

This is the preferred method for the Agent-based installation. Using GitOps ZTP manifests is optional.

- Create the **install-config.yaml** file:

```
$ cat << EOF > ./my-cluster/install-config.yaml
apiVersion: v1
baseDomain: test.example.com
compute:
- architecture: amd64 1
  hyperthreading: Enabled
  name: worker
  replicas: 0
controlPlane:
  architecture: amd64
  hyperthreading: Enabled
  name: master
  replicas: 1
metadata:
  name: sno-cluster 2
networking:
  clusterNetwork:
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  machineNetwork:
  - cidr: 192.168.0.0/16
  networkType: OVNKubernetes 3
  serviceNetwork:
  - 172.30.0.0/16
platform: 4
  none: {}
pullSecret: '<pull_secret>' 5
sshKey: '<ssh_pub_key>' 6
EOF
```

- Specify the system architecture. Valid values are **amd64**, **arm64**, **ppc64le**, and **s390x**.

If you are using the release image with the **multi** payload, you can install the cluster on different architectures such as **arm64**, **amd64**, **s390x**, and **ppc64le**. Otherwise, you can install the cluster only on the **release architecture** displayed in the output of the **openshift-install version** command. For more information, see "Verifying the supported architecture for installing an Agent-based Installer cluster".

- 2 Required. Specify your cluster name.
- 3 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 4 Specify your platform.

**NOTE**

For bare metal platforms, host settings made in the platform section of the **install-config.yaml** file are used by default, unless they are overridden by configurations made in the **agent-config.yaml** file.

- 5 Specify your pull secret.
- 6 Specify your SSH public key.

**NOTE**

If you set the platform to **vSphere** or **baremetal**, you can configure IP address endpoints for cluster nodes in three ways:

- IPv4
- IPv6
- IPv4 and IPv6 in parallel (dual-stack)

IPv6 is supported only on bare metal platforms.

Example of dual-stack networking

```
networking:
  clusterNetwork:
    - cidr: 172.21.0.0/16
      hostPrefix: 23
    - cidr: fd02::/48
      hostPrefix: 64
  machineNetwork:
    - cidr: 192.168.11.0/16
    - cidr: 2001:DB8::/32
  serviceNetwork:
    - 172.22.0.0/16
    - fd03::/112
  networkType: OVNKubernetes
platform:
  baremetal:
    apiVIPs:
      - 192.168.11.3
      - 2001:DB8::4
    ingressVIPs:
      - 192.168.11.4
      - 2001:DB8::5
```

**NOTE**

When you use a disconnected mirror registry, you must add the certificate file that you created previously for your mirror registry to the **additionalTrustBundle** field of the **install-config.yaml** file.

5. Create the **agent-config.yaml** file:

```
$ cat > agent-config.yaml << EOF
apiVersion: v1beta1
kind: AgentConfig
metadata:
  name: sno-cluster
rendezvousIP: 192.168.111.80 1
hosts: 2
  - hostname: master-0 3
  interfaces:
    - name: eno1
      macAddress: 00:ef:44:21:e6:a5
  rootDeviceHints: 4
    deviceName: /dev/sdb
  networkConfig: 5
    interfaces:
      - name: eno1
        type: ethernet
        state: up
        mac-address: 00:ef:44:21:e6:a5
        ipv4:
          enabled: true
          address:
            - ip: 192.168.111.80
              prefix-length: 23
          dhcp: false
    dns-resolver:
      config:
        server:
          - 192.168.111.1
    routes:
      config:
        - destination: 0.0.0.0/0
          next-hop-address: 192.168.111.2
          next-hop-interface: eno1
          table-id: 254
EOF
```

- 1** This IP address is used to determine which node performs the bootstrapping process as well as running the **assisted-service** component. You must provide the rendezvous IP address when you do not specify at least one host's IP address in the **networkConfig** parameter. If this address is not provided, one IP address is selected from the provided hosts' **networkConfig**.
- 2** Optional: Host configuration. The number of hosts defined must not exceed the total number of hosts defined in the **install-config.yaml** file, which is the sum of the values of the **compute.replicas** and **controlPlane.replicas** parameters.

- 3 Optional: Overrides the hostname obtained from either the Dynamic Host Configuration Protocol (DHCP) or a reverse DNS lookup. Each host must have a unique hostname
- 4 Enables provisioning of the Red Hat Enterprise Linux CoreOS (RHCOS) image to a particular device. The installation program examines the devices in the order it discovers them, and compares the discovered values with the hint values. It uses the first discovered device that matches the hint value.
- 5 Optional: Configures the network interface of a host in NMState format.

Additional resources

- [Configuring regions and zones for a VMware vCenter](#)
- [Verifying the supported architecture for installing an Agent-based installer cluster](#)
- [Configuring the Agent-based Installer to use mirrored images](#)

13.3.2.4. Optional: Creating additional manifest files

You can create additional manifests to further configure your cluster beyond the configurations available in the **install-config.yaml** and **agent-config.yaml** files.

13.3.2.4.1. Creating a directory to contain additional manifests

If you create additional manifests to configure your Agent-based installation beyond the **install-config.yaml** and **agent-config.yaml** files, you must create an **openshift** subdirectory within your installation directory. All of your additional machine configurations must be located within this subdirectory.



NOTE

The most common type of additional manifest you can add is a **MachineConfig** object. For examples of **MachineConfig** objects you can add during the Agent-based installation, see "Using MachineConfig objects to configure nodes" in the "Additional resources" section.

Procedure

- On your installation host, create an **openshift** subdirectory within the installation directory by running the following command:

```
$ mkdir <installation_directory>/openshift
```

Additional resources

- [Using MachineConfig objects to configure nodes](#)

13.3.2.4.2. Disk partitioning

In general, you should use the default disk partitioning that is created during the RHCOS installation. However, there are cases where you might want to create a separate partition for a directory that you expect to grow.

OpenShift Container Platform supports the addition of a single partition to attach storage to either the `/var` directory or a subdirectory of `/var`. For example:

- `/var/lib/containers`: Holds container-related content that can grow as more images and containers are added to a system.
- `/var/lib/etcd`: Holds data that you might want to keep separate for purposes such as performance optimization of etcd storage.
- `/var`: Holds data that you might want to keep separate for purposes such as auditing.



IMPORTANT

For disk sizes larger than 100GB, and especially larger than 1TB, create a separate `/var` partition.

Storing the contents of a `/var` directory separately makes it easier to grow storage for those areas as needed and reinstall OpenShift Container Platform at a later date and keep that data intact. With this method, you will not have to pull all your containers again, nor will you have to copy massive log files when you update systems.

The use of a separate partition for the `/var` directory or a subdirectory of `/var` also prevents data growth in the partitioned directory from filling up the root file system.

The following procedure sets up a separate `/var` partition by adding a machine config manifest that is wrapped into the Ignition config file for a node type during the preparation phase of an installation.

Prerequisites

- You have created an `openshift` subdirectory within your installation directory.

Procedure

1. Create a Butane config that configures the additional partition. For example, name the file `$HOME/clusterconfig/98-var-partition.bu`, change the disk device name to the name of the storage device on the `worker` systems, and set the storage size as appropriate. This example places the `/var` directory on a separate partition:

```
variant: openshift
version: 4.16.0
metadata:
  labels:
    machineconfiguration.openshift.io/role: worker
  name: 98-var-partition
storage:
  disks:
    - device: /dev/disk/by-id/<device_name> 1
    partitions:
      - label: var
        start_mib: <partition_start_offset> 2
        size_mib: <partition_size> 3
        number: 5
  filesystems:
    - device: /dev/disk/by-partlabel/var
      path: /var
```

```
format: xfs
mount_options: [defaults, prjquota] 4
with_mount_unit: true
```

- 1 The storage device name of the disk that you want to partition.
- 2 When adding a data partition to the boot disk, a minimum offset value of 25000 mebibytes is recommended. The root file system is automatically resized to fill all available space up to the specified offset. If no offset value is specified, or if the specified value is smaller than the recommended minimum, the resulting root file system will be too small, and future reinstalls of RHCOS might overwrite the beginning of the data partition.
- 3 The size of the data partition in mebibytes.
- 4 The **prjquota** mount option must be enabled for filesystems used for container storage.



NOTE

When creating a separate **/var** partition, you cannot use different instance types for compute nodes, if the different instance types do not have the same device name.

2. Create a manifest from the Butane config and save it to the **clusterconfig/openshift** directory. For example, run the following command:

```
$ butane $HOME/clusterconfig/98-var-partition.bu -o $HOME/clusterconfig/openshift/98-var-partition.yaml
```

13.3.2.5. Optional: Using ZTP manifests

You can use GitOps Zero Touch Provisioning (ZTP) manifests to configure your installation beyond the options available through the **install-config.yaml** and **agent-config.yaml** files.



NOTE

GitOps ZTP manifests can be generated with or without configuring the **install-config.yaml** and **agent-config.yaml** files beforehand. If you chose to configure the **install-config.yaml** and **agent-config.yaml** files, the configurations will be imported to the ZTP cluster manifests when they are generated.

Prerequisites

- You have placed the **openshift-install** binary in a directory that is on your **PATH**.
- Optional: You have created and configured the **install-config.yaml** and **agent-config.yaml** files.

Procedure

1. Use the following command to generate ZTP cluster manifests:

```
$ openshift-install agent create cluster-manifests --dir <installation_directory>
```



IMPORTANT

If you have created the **install-config.yaml** and **agent-config.yaml** files, those files are deleted and replaced by the cluster manifests generated through this command.

Any configurations made to the **install-config.yaml** and **agent-config.yaml** files are imported to the ZTP cluster manifests when you run the **openshift-install agent create cluster-manifests** command.

2. Navigate to the **cluster-manifests** directory:

```
$ cd <installation_directory>/cluster-manifests
```

3. Configure the manifest files in the **cluster-manifests** directory. For sample files, see the "Sample GitOps ZTP custom resources" section.
4. Disconnected clusters: If you did not define mirror configuration in the **install-config.yaml** file before generating the ZTP manifests, perform the following steps:
 - a. Navigate to the **mirror** directory:

```
$ cd ../mirror
```

- b. Configure the manifest files in the **mirror** directory.

Additional resources

- [Sample GitOps ZTP custom resources](#) .
- See [Challenges of the network far edge](#) to learn more about GitOps Zero Touch Provisioning (ZTP).

13.3.2.6. Optional: Encrypting the disk

Use this procedure to encrypt your disk or partition while installing OpenShift Container Platform with the Agent-based Installer.

Prerequisites

- You have created and configured the **install-config.yaml** and **agent-config.yaml** files, unless you are using ZTP manifests.
- You have placed the **openshift-install** binary in a directory that is on your **PATH**.

Procedure

1. Use the following command to generate ZTP cluster manifests:

```
$ openshift-install agent create cluster-manifests --dir <installation_directory>
```




IMPORTANT

If you have created the **install-config.yaml** and **agent-config.yaml** files, those files are deleted and replaced by the cluster manifests generated through this command.

Any configurations made to the **install-config.yaml** and **agent-config.yaml** files are imported to the ZTP cluster manifests when you run the **openshift-install agent create cluster-manifests** command.



NOTE

If you have already generated ZTP manifests, skip this step.

2. Navigate to the **cluster-manifests** directory:

```
$ cd <installation_directory>/cluster-manifests
```

3. Add the following section to the **agent-cluster-install.yaml** file:

```
diskEncryption:
  enableOn: all 1
  mode: tang 2
  tangServers: "server1": "http://tang-server-1.example.com:7500" 3
```

- 1** Specify which nodes to enable disk encryption on. Valid values are 'none', 'all', 'master', and 'worker'.
- 2** Specify which disk encryption mode to use. Valid values are 'tpmv2' and 'tang'.
- 3** Optional: If you are using Tang, specify the Tang servers.

Additional resources

- [About disk encryption](#)

13.3.2.7. Creating and booting the agent image

Use this procedure to boot the agent image on your machines.

Procedure

1. Create the agent image by running the following command:

```
$ openshift-install --dir <install_directory> agent create image
```



NOTE

Red Hat Enterprise Linux CoreOS (RHCOS) supports multipathing on the primary disk, allowing stronger resilience to hardware failure to achieve higher host availability. Multipathing is enabled by default in the agent ISO image, with a default **/etc/multipath.conf** configuration.

2. Boot the **agent.x86_64.iso**, **agent.aarch64.iso**, or **agent.s390x.iso** image on the bare metal machines.

13.3.2.8. Adding IBM Z(R) agents with RHEL KVM

Use the following procedure to manually add IBM Z® agents with RHEL KVM.



NOTE

Currently, ISO boot support on IBM Z® (**s390x**) is available only for RHEL KVM, which provides the flexibility to choose either PXE or ISO-based installation. For installations with z/VM, only PXE boot is supported.

Procedure

1. Boot your RHEL KVM machine.
2. To deploy the virtual server, run the **virt-install** command with the following parameters:

```
$ virt-install
  --name <vm_name> \
  --autostart \
  --memory=<memory> \
  --cpu host \
  --vcpus=<vcpus> \
  --cdrom <agent.iso_image> \ 1
  --disk pool=default,size=<disk_pool_size> \
  --network network:default,mac=<mac_address> \
  --graphics none \
  --noautoconsole \
  --os-variant rhel9.0 \
  --wait=-1
```



For the **--cdrom** parameter, specify the location of the ISO image on the HTTP or HTTPS server.

13.3.2.9. Verifying that the current installation host can pull release images

After you boot the agent image and network services are made available to the host, the agent console application performs a pull check to verify that the current host can retrieve release images.

If the primary pull check passes, you can quit the application to continue with the installation. If the pull check fails, the application performs additional checks, as seen in the **Additional checks** section of the TUI, to help you troubleshoot the problem. A failure for any of the additional checks is not necessarily critical as long as the primary pull check succeeds.

If there are host network configuration issues that might cause an installation to fail, you can use the console application to make adjustments to your network configurations.



IMPORTANT

If the agent console application detects host network configuration issues, the installation workflow will be halted until the user manually stops the console application and signals the intention to proceed.

Procedure

1. Wait for the agent console application to check whether or not the configured release image can be pulled from a registry.
2. If the agent console application states that the installer connectivity checks have passed, wait for the prompt to time out to continue with the installation.



NOTE

You can still choose to view or change network configuration settings even if the connectivity checks have passed.

However, if you choose to interact with the agent console application rather than letting it time out, you must manually quit the TUI to proceed with the installation.

3. If the agent console application checks have failed, which is indicated by a red icon beside the **Release image URL** pull check, use the following steps to reconfigure the host's network settings:
 - a. Read the **Check Errors** section of the TUI. This section displays error messages specific to the failed checks.

```

Agent installer network boot setup
-----
Release image URL
-----
■ quay.io/openshift-release-dev/ocp-release:4.13.0-x86_64

Additional checks
-----
■ nslookup quay.io
■ ping quay.io
■ quay.io responds to http GET

Check Errors
-----
ping: quay.io: Name or service not known
ping failure:
ping: quay.io: Name or service not known
nslookup failure:
;; connection timed out; no servers could be reached

ping failure:
ping: quay.io: Name or service not known

<Configure network> <Quit>
  
```

- b. Select **Configure network** to launch the NetworkManager TUI.

- c. Select **Edit a connection** and select the connection you want to reconfigure.
- d. Edit the configuration and select **OK** to save your changes.
- e. Select **Back** to return to the main screen of the NetworkManager TUI.
- f. Select **Activate a Connection**.
- g. Select the reconfigured network to deactivate it.
- h. Select the reconfigured network again to reactivate it.
- i. Select **Back** and then select **Quit** to return to the agent console application.
- j. Wait at least five seconds for the continuous network checks to restart using the new network configuration.
- k. If the **Release image URL** pull check succeeds and displays a green icon beside the URL, select **Quit** to exit the agent console application and continue with the installation.

13.3.2.10. Tracking and verifying installation progress

Use the following procedure to track installation progress and to verify a successful installation.

Prerequisites

- You have configured a DNS record for the Kubernetes API server.

Procedure

1. Optional: To know when the bootstrap host (rendezvous host) reboots, run the following command:

```
$ ./openshift-install --dir <install_directory> agent wait-for bootstrap-complete \1
--log-level=info 2
```

- 1 For **<install_directory>**, specify the path to the directory where the agent ISO was generated.

- 2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Example output

```
.....
.....
INFO Bootstrap configMap status is complete
INFO cluster bootstrap is complete
```

The command succeeds when the Kubernetes API server signals that it has been bootstrapped on the control plane machines.

2. To track the progress and verify successful installation, run the following command:

```
$ openshift-install --dir <install_directory> agent wait-for install-complete 1
```

- 1 For `<install_directory>` directory, specify the path to the directory where the agent ISO was generated.

Example output

```

.....
.....
INFO Cluster is installed
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run
INFO   export KUBECONFIG=/home/core/installer/auth/kubeconfig
INFO Access the OpenShift web-console here: https://console-openshift-console.apps.sno-
cluster.test.example.com

```

NOTE

If you are using the optional method of GitOps ZTP manifests, you can configure IP address endpoints for cluster nodes through the **AgentClusterInstall.yaml** file in three ways:

- IPv4
- IPv6
- IPv4 and IPv6 in parallel (dual-stack)

IPv6 is supported only on bare metal platforms.

Example of dual-stack networking

```

apiVIP: 192.168.11.3
ingressVIP: 192.168.11.4
clusterDeploymentRef:
  name: mycluster
imageSetRef:
  name: openshift-4.16
networking:
  clusterNetwork:
    - cidr: 172.21.0.0/16
      hostPrefix: 23
    - cidr: fd02::/48
      hostPrefix: 64
  machineNetwork:
    - cidr: 192.168.11.0/16
    - cidr: 2001:DB8::/32
  serviceNetwork:
    - 172.22.0.0/16
    - fd03::/112
  networkType: OVNKubernetes

```

Additional resources

- See [Deploying with dual-stack networking](#) .
- See [Configuring the install-config yaml file](#).
- See [Configuring a three-node cluster](#) to deploy three-node clusters in bare metal environments.
- See [About root device hints](#) .
- See [NMState state examples](#).

13.3.3. Sample GitOps ZTP custom resources

Optional: You can use GitOps Zero Touch Provisioning (ZTP) custom resource (CR) objects to install an OpenShift Container Platform cluster with the Agent-based Installer.

You can customize the following GitOps ZTP custom resources to specify more details about your OpenShift Container Platform cluster. The following sample GitOps ZTP custom resources are for a single-node cluster.

agent-cluster-install.yaml

```
apiVersion: extensions.hive.openshift.io/v1beta1
kind: AgentClusterInstall
metadata:
  name: test-agent-cluster-install
  namespace: cluster0
spec:
  clusterDeploymentRef:
    name: ostest
  imageSetRef:
    name: openshift-4.16
  networking:
    clusterNetwork:
      - cidr: 10.128.0.0/14
        hostPrefix: 23
    serviceNetwork:
      - 172.30.0.0/16
  provisionRequirements:
    controlPlaneAgents: 1
    workerAgents: 0
  sshPublicKey: <YOUR_SSH_PUBLIC_KEY>
```

cluster-deployment.yaml

```
apiVersion: hive.openshift.io/v1
kind: ClusterDeployment
metadata:
  name: ostest
  namespace: cluster0
spec:
  baseDomain: test.metalkube.org
  clusterInstallRef:
    group: extensions.hive.openshift.io
    kind: AgentClusterInstall
```

```

name: test-agent-cluster-install
version: v1beta1
clusterName: ostest
controlPlaneConfig:
  servingCertificates: {}
platform:
  agentBareMetal:
    agentSelector:
      matchLabels:
        bla: aaa
pullSecretRef:
  name: pull-secret

```

cluster-image-set.yaml

```

apiVersion: hive.openshift.io/v1
kind: ClusterImageSet
metadata:
  name: openshift-4.16
spec:
  releaseImage: registry.ci.openshift.org/ocp/release:4.16.0-0.nightly-2022-06-06-025509

```

infra-env.yaml

```

apiVersion: agent-install.openshift.io/v1beta1
kind: InfraEnv
metadata:
  name: myinfraenv
  namespace: cluster0
spec:
  clusterRef:
    name: ostest
    namespace: cluster0
  cpuArchitecture: aarch64
  pullSecretRef:
    name: pull-secret
  sshAuthorizedKey: <YOUR_SSH_PUBLIC_KEY>
  nmStateConfigLabelSelector:
    matchLabels:
      cluster0-nmstate-label-name: cluster0-nmstate-label-value

```

nmstateconfig.yaml

```

apiVersion: agent-install.openshift.io/v1beta1
kind: NMStateConfig
metadata:
  name: master-0
  namespace: openshift-machine-api
  labels:
    cluster0-nmstate-label-name: cluster0-nmstate-label-value
spec:
  config:
    interfaces:
      - name: eth0
        type: ethernet

```

```

state: up
mac-address: 52:54:01:aa:aa:a1
ipv4:
  enabled: true
  address:
    - ip: 192.168.122.2
      prefix-length: 23
  dhcp: false
dns-resolver:
  config:
    server:
      - 192.168.122.1
routes:
  config:
    - destination: 0.0.0.0/0
      next-hop-address: 192.168.122.1
      next-hop-interface: eth0
      table-id: 254
interfaces:
  - name: "eth0"
    macAddress: 52:54:01:aa:aa:a1

```

pull-secret.yaml

```

apiVersion: v1
kind: Secret
type: kubernetes.io/dockerconfigjson
metadata:
  name: pull-secret
  namespace: cluster0
stringData:
  .dockerconfigjson: 'YOUR_PULL_SECRET'

```

Additional resources

- See [Challenges of the network far edge](#) to learn more about GitOps Zero Touch Provisioning (ZTP).

13.3.4. Gathering log data from a failed Agent-based installation

Use the following procedure to gather log data about a failed Agent-based installation to provide for a support case.

Prerequisites

- You have configured a DNS record for the Kubernetes API server.

Procedure

1. Run the following command and collect the output:

```

$ ./openshift-install --dir <install_directory> agent wait-for bootstrap-complete --log-level=debug

```


Example error message

```
...
ERROR Bootstrap failed to complete: : bootstrap process timed out: context deadline
exceeded
```

- If the output from the previous command indicates a failure, or if the bootstrap is not progressing, run the following command to connect to the rendezvous host and collect the output:

```
$ ssh core@<node-ip> agent-gather -O >agent-gather.tar.xz
```



NOTE

Red Hat Support can diagnose most issues using the data gathered from the rendezvous host, but if some hosts are not able to register, gathering this data from every host might be helpful.

- If the bootstrap completes and the cluster nodes reboot, run the following command and collect the output:

```
$ ./openshift-install --dir <install_directory> agent wait-for install-complete --log-level=debug
```

- If the output from the previous command indicates a failure, perform the following steps:
 - Export the **kubeconfig** file to your environment by running the following command:

```
$ export KUBECONFIG=<install_directory>/auth/kubeconfig
```

- To gather information for debugging, run the following command:

```
$ oc adm must-gather
```

- Create a compressed file from the **must-gather** directory that was just created in your working directory by running the following command:

```
$ tar cvaf must-gather.tar.gz <must_gather_directory>
```

- Excluding the **/auth** subdirectory, attach the installation directory used during the deployment to your support case on the [Red Hat Customer Portal](#).
- Attach all other data gathered from this procedure to your support case.

13.4. PREPARING PXE ASSETS FOR OPENSIFT CONTAINER PLATFORM

Use the following procedures to create the assets needed to PXE boot an OpenShift Container Platform cluster using the Agent-based Installer.

The assets you create in these procedures will deploy a single-node OpenShift Container Platform installation. You can use these procedures as a basis and modify configurations according to your requirements.

13.4.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.

13.4.2. Downloading the Agent-based Installer

Use this procedure to download the Agent-based Installer and the CLI needed for your installation.

Procedure

1. Log in to the OpenShift Container Platform web console using your login credentials.
2. Navigate to [Datacenter](#).
3. Click **Run Agent-based Installer locally**.
4. Select the operating system and architecture for the **OpenShift Installer** and **Command line interface**.
5. Click **Download Installer** to download and extract the install program.
6. You can either download or copy the pull secret by clicking on **Download pull secret** or **Copy pull secret**.
7. Click **Download command-line tools** and place the **openshift-install** binary in a directory that is on your **PATH**.

13.4.3. Creating the preferred configuration inputs

Use this procedure to create the preferred configuration inputs used to create the PXE files.

Procedure

1. Install **nmstate** dependency by running the following command:

```
$ sudo dnf install /usr/bin/nmstatectl -y
```

2. Place the **openshift-install** binary in a directory that is on your PATH.
3. Create a directory to store the install configuration by running the following command:

```
$ mkdir ~/<directory_name>
```



NOTE

This is the preferred method for the Agent-based installation. Using GitOps ZTP manifests is optional.

4. Create the **install-config.yaml** file:

```
$ cat << EOF > ./my-cluster/install-config.yaml
apiVersion: v1
baseDomain: test.example.com
```

```

compute:
- architecture: amd64 1
  hyperthreading: Enabled
  name: worker
  replicas: 0
controlPlane:
  architecture: amd64
  hyperthreading: Enabled
  name: master
  replicas: 1
metadata:
  name: sno-cluster 2
networking:
  clusterNetwork:
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  machineNetwork:
  - cidr: 192.168.0.0/16
  networkType: OVNKubernetes 3
  serviceNetwork:
  - 172.30.0.0/16
platform: 4
  none: {}
pullSecret: '<pull_secret>' 5
sshKey: '<ssh_pub_key>' 6
EOF

```

- 1** Specify the system architecture. Valid values are **amd64**, **arm64**, **ppc64le**, and **s390x**.

If you are using the release image with the **multi** payload, you can install the cluster on different architectures such as **arm64**, **amd64**, **s390x**, and **ppc64le**. Otherwise, you can install the cluster only on the **release architecture** displayed in the output of the **openshift-install version** command. For more information, see "Verifying the supported architecture for installing an Agent-based Installer cluster".

- 2** Required. Specify your cluster name.
- 3** The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 4** Specify your platform.



NOTE

For bare metal platforms, host settings made in the platform section of the **install-config.yaml** file are used by default, unless they are overridden by configurations made in the **agent-config.yaml** file.

- 5** Specify your pull secret.
- 6** Specify your SSH public key.

**NOTE**

If you set the platform to **vSphere** or **baremetal**, you can configure IP address endpoints for cluster nodes in three ways:

- IPv4
- IPv6
- IPv4 and IPv6 in parallel (dual-stack)

IPv6 is supported only on bare metal platforms.

Example of dual-stack networking

```
networking:
  clusterNetwork:
    - cidr: 172.21.0.0/16
      hostPrefix: 23
    - cidr: fd02::/48
      hostPrefix: 64
  machineNetwork:
    - cidr: 192.168.11.0/16
    - cidr: 2001:DB8::/32
  serviceNetwork:
    - 172.22.0.0/16
    - fd03::/112
  networkType: OVNKubernetes
platform:
  baremetal:
    apiVIPs:
      - 192.168.11.3
      - 2001:DB8::4
    ingressVIPs:
      - 192.168.11.4
      - 2001:DB8::5
```

**NOTE**

When you use a disconnected mirror registry, you must add the certificate file that you created previously for your mirror registry to the **additionalTrustBundle** field of the **install-config.yaml** file.

5. Create the **agent-config.yaml** file:

```
$ cat > agent-config.yaml << EOF
apiVersion: v1beta1
kind: AgentConfig
metadata:
  name: sno-cluster
rendezvousIP: 192.168.111.80 1
hosts: 2
  - hostname: master-0 3
interfaces:
```

```

- name: eno1
  macAddress: 00:ef:44:21:e6:a5
rootDeviceHints: 4
  deviceName: /dev/sdb
networkConfig: 5
  interfaces:
    - name: eno1
      type: ethernet
      state: up
      mac-address: 00:ef:44:21:e6:a5
      ipv4:
        enabled: true
        address:
          - ip: 192.168.111.80
            prefix-length: 23
        dhcp: false
  dns-resolver:
    config:
      server:
        - 192.168.111.1
  routes:
    config:
      - destination: 0.0.0.0/0
        next-hop-address: 192.168.111.2
        next-hop-interface: eno1
        table-id: 254
EOF

```

- 1 This IP address is used to determine which node performs the bootstrapping process as well as running the **assisted-service** component. You must provide the rendezvous IP address when you do not specify at least one host's IP address in the **networkConfig** parameter. If this address is not provided, one IP address is selected from the provided hosts' **networkConfig**.
- 2 Optional: Host configuration. The number of hosts defined must not exceed the total number of hosts defined in the **install-config.yaml** file, which is the sum of the values of the **compute.replicas** and **controlPlane.replicas** parameters.
- 3 Optional: Overrides the hostname obtained from either the Dynamic Host Configuration Protocol (DHCP) or a reverse DNS lookup. Each host must have a unique hostname supplied by one of these methods.
- 4 Enables provisioning of the Red Hat Enterprise Linux CoreOS (RHCOS) image to a particular device. The installation program examines the devices in the order it discovers them, and compares the discovered values with the hint values. It uses the first discovered device that matches the hint value.
- 5 Optional: Configures the network interface of a host in NMState format.

6. Optional: To create an iPXE script, add the **bootArtifactsBaseURL** to the **agent-config.yaml** file:

```

apiVersion: v1beta1
kind: AgentConfig
metadata:

```

```

name: sno-cluster
rendezvousIP: 192.168.111.80
bootArtifactsBaseURL: <asset_server_URL>

```

Where **<asset_server_URL>** is the URL of the server you will upload the PXE assets to.

Additional resources

- [Deploying with dual-stack networking](#) .
- [Configuring the install-config yaml file](#) .
- See [Configuring a three-node cluster](#) to deploy three-node clusters in bare metal environments.
- [About root device hints](#) .
- [NMState state examples](#) .
- [Optional: Creating additional manifest files](#)

13.4.4. Creating the PXE assets

Use the following procedure to create the assets and optional script to implement in your PXE infrastructure.

Procedure

1. Create the PXE assets by running the following command:

```
$ openshift-install agent create pxe-files
```


The generated PXE assets and optional iPXE script can be found in the **boot-artifacts** directory.

Example filesystem with PXE assets and optional iPXE script

```

boot-artifacts
├── agent.x86_64-initrd.img
├── agent.x86_64.ipxe
├── agent.x86_64-rootfs.img
└── agent.x86_64-vmlinuz

```



IMPORTANT

The contents of the **boot-artifacts** directory vary depending on the specified architecture.



NOTE

Red Hat Enterprise Linux CoreOS (RHCOS) supports multipathing on the primary disk, allowing stronger resilience to hardware failure to achieve higher host availability. Multipathing is enabled by default in the agent ISO image, with a default **/etc/multipath.conf** configuration.

2. Upload the PXE assets and optional script to your infrastructure where they will be accessible during the boot process.



NOTE

If you generated an iPXE script, the location of the assets must match the **bootArtifactsBaseURL** you added to the **agent-config.yaml** file.

13.4.5. Manually adding IBM Z agents

After creating the PXE assets, you can add IBM Z[®] agents.

Depending on your IBM Z[®] environment, you can choose from the following options:

- Adding IBM Z[®] agents with z/VM
- Adding IBM Z[®] agents with RHEL KVM

13.4.5.1. Adding IBM Z agents with z/VM

Use the following procedure to manually add IBM Z[®] agents with z/VM.

Procedure

1. Create a parameter file for the z/VM guest:

Example parameter file

```
rd.neednet=1 \
console=ttysclp0 \
coreos.live.rootfs_url=<rootfs_url> \ 1
ip=172.18.78.2::172.18.78.1:255.255.255.0:::none nameserver=172.18.78.1 \ 2
zfcf.allow_lun_scan=0 \ 3
rd.znet=qeth,0.0.bdd0,0.0.bdd1,0.0.bdd2,layer2=1 \
rd.dasd=0.0.4411 \ 4
rd.zfcf=0.0.8001,0x50050763040051e3,0x4000406300000000 \ 5
random.trust_cpu=on rd.luks.options=discard \
ignition.firstboot ignition.platform.id=metal \
console=ty1 console=tyS1,115200n8 \
coreos.inst.persistent-kargs="console=ty1 console=tyS1,115200n8"
```

- 1 For the **coreos.live.rootfs_url** artifact, specify the matching **rootfs** artifact for the **kernel** and **initramfs** that you are booting. Only HTTP and HTTPS protocols are supported.
- 2 For the **ip** parameter, assign the IP address automatically using DHCP, or manually assign the IP address, as described in "Installing a cluster with z/VM on IBM Z[®] and IBM[®] LinuxONE".
- 3 The default is **1**. Omit this entry when using an OSA network adapter.
- 4 For installations on DASD-type disks, use **rd.dasd** to specify the DASD where Red Hat Enterprise Linux CoreOS (RHCOS) is to be installed. Omit this entry for FCP-type disks.
- 5 For installations on FCP-type disks, use **rd.zfcf=<adapter>,<wwpn>,<lun>** to specify the FCP disk where RHCOS is to be installed. Omit this entry for DASD-type disks.

to the disk where RHECOS is to be installed. Omit this entry for `BR/SD` type disks.

Leave all other parameters unchanged.

2. Punch the **kernel.img.generic.parm**, and **initrd.img** files to the virtual reader of the z/VM guest virtual machine.

For more information, see [PUNCH](#) in IBM Documentation.

TIP

You can use the **CP PUNCH** command or, if you use Linux, the **vmur** command, to transfer files between two z/VM guest virtual machines.

3. Log in to the conversational monitor system (CMS) on the bootstrap machine.
4. IPL the bootstrap machine from the reader by running the following command:

```
$ ipl c
```

For more information, see [IPL](#) in IBM Documentation.

13.4.5.2. Adding IBM Z(R) agents with RHEL KVM

Use the following procedure to manually add IBM Z® agents with RHEL KVM.



NOTE

Currently, ISO boot support on IBM Z® (**s390x**) is available only for RHEL KVM, which provides the flexibility to choose either PXE or ISO-based installation. For installations with z/VM, only PXE boot is supported.

Procedure

1. Boot your RHEL KVM machine.
2. To deploy the virtual server, run the **virt-install** command with the following parameters:

```
$ virt-install \
  --name <vm_name> \
  --autostart \
  --ram=16384 \
  --cpu host \
  --vcpus=8 \
  --location <path_to_kernel_initrd_image>,kernel=kernel.img,initrd=initrd.img \ 1
  --disk <qcow_image_path> \
  --network network:macvtap ,mac=<mac_address> \
  --graphics none \
  --noautoconsole \
  --wait=-1 \
  --extra-args "rd.neednet=1 nameserver=<nameserver>" \
  --extra-args "ip=<IP>:::<nameserver>:::<hostname>:enc1:none" \
  --extra-args "coreos.live.rootfs_url=http://<http_server>:8080/agent.s390x-rootfs.img" \
  --extra-args "random.trust_cpu=on rd.luks.options=discard" \
  --extra-args "ignition.firstboot ignition.platform.id=metal" \
```



```
--extra-args "console=tty1 console=ttyS1,115200n8" \
--extra-args "coreos.inst.persistent-kargs=console=tty1 console=ttyS1,115200n8" \
--osinfo detect=on,require=off
```

- 1 For the **--location** parameter, specify the location of the kernel/initrd on the HTTP or HTTPS server.

13.4.6. Additional resources

- See [Installing an OpenShift Container Platform cluster with the Agent-based Installer](#) to learn about more configurations available with the Agent-based Installer.

13.5. PREPARING AN AGENT-BASED INSTALLED CLUSTER FOR THE MULTICLUSTER ENGINE FOR KUBERNETES OPERATOR

You can install the multicluster engine Operator and deploy a hub cluster with the Agent-based OpenShift Container Platform Installer. The following procedure is partially automated and requires manual steps after the initial cluster is deployed.

13.5.1. Prerequisites

- You have read the following documentation:
 - [Cluster lifecycle with multicluster engine operator overview](#) .
 - [Persistent storage using local volumes](#) .
 - [Using GitOps ZTP to provision clusters at the network far edge](#) .
 - [Preparing to install with the Agent-based Installer](#) .
 - [About disconnected installation mirroring](#) .
- You have access to the internet to obtain the necessary container images.
- You have installed the OpenShift CLI (**oc**).
- If you are installing in a disconnected environment, you must have a configured local mirror registry for disconnected installation mirroring.

13.5.2. Preparing an Agent-based cluster deployment for the multicluster engine for Kubernetes Operator while disconnected

You can mirror the required OpenShift Container Platform container images, the multicluster engine Operator, and the Local Storage Operator (LSO) into your local mirror registry in a disconnected environment. Ensure that you note the local DNS hostname and port of your mirror registry.



NOTE

To mirror your OpenShift Container Platform image repository to your mirror registry, you can use either the **oc adm release image** or **oc mirror** command. In this procedure, the **oc mirror** command is used as an example.

Procedure

1. Create an **<assets_directory>** folder to contain valid **install-config.yaml** and **agent-config.yaml** files. This directory is used to store all the assets.
2. To mirror an OpenShift Container Platform image repository, the multicluster engine, and the LSO, create a **ImageSetConfiguration.yaml** file with the following settings:

Example ImageSetConfiguration.yaml

```
kind: ImageSetConfiguration
apiVersion: mirror.openshift.io/v1alpha2
archiveSize: 4 1
storageConfig: 2
  imageURL: <your-local-registry-dns-name>:<your-local-registry-port>/mirror/oc-mirror-
  metadata 3
    skipTLS: true
  mirror:
    platform:
      architectures:
        - "amd64"
      channels:
        - name: stable-4.16 4
          type: ocp
    additionalImages:
      - name: registry.redhat.io/ubi9/ubi:latest
    operators:
      - catalog: registry.redhat.io/redhat/redhat-operator-index:v4.16 5
        packages: 6
          - name: multicluster-engine 7
          - name: local-storage-operator 8
```

- 1** Specify the maximum size, in GiB, of each file within the image set.
- 2** Set the back-end location to receive the image set metadata. This location can be a registry or local directory. It is required to specify **storageConfig** values.
- 3** Set the registry URL for the storage backend.
- 4** Set the channel that contains the OpenShift Container Platform images for the version you are installing.
- 5** Set the Operator catalog that contains the OpenShift Container Platform images that you are installing.
- 6** Specify only certain Operator packages and channels to include in the image set. Remove this field to retrieve all packages in the catalog.
- 7** The multicluster engine packages and channels.
- 8** The LSO packages and channels.

**NOTE**

This file is required by the **oc mirror** command when mirroring content.

- To mirror a specific OpenShift Container Platform image repository, the multicluster engine, and the LSO, run the following command:

```
$ oc mirror --dest-skip-tls --config ocp-mce-imageset.yaml docker://<your-local-registry-dns-name>:<your-local-registry-port>
```

- Update the registry and certificate in the **install-config.yaml** file:


Example imageContentSources.yaml

```
imageContentSources:
- source: "quay.io/openshift-release-dev/ocp-release"
  mirrors:
  - "<your-local-registry-dns-name>:<your-local-registry-port>/openshift/release-images"
- source: "quay.io/openshift-release-dev/ocp-v4.0-art-dev"
  mirrors:
  - "<your-local-registry-dns-name>:<your-local-registry-port>/openshift/release"
- source: "registry.redhat.io/ubi9"
  mirrors:
  - "<your-local-registry-dns-name>:<your-local-registry-port>/ubi9"
- source: "registry.redhat.io/multicluster-engine"
  mirrors:
  - "<your-local-registry-dns-name>:<your-local-registry-port>/multicluster-engine"
- source: "registry.redhat.io/rhel8"
  mirrors:
  - "<your-local-registry-dns-name>:<your-local-registry-port>/rhel8"
- source: "registry.redhat.io/redhat"
  mirrors:
  - "<your-local-registry-dns-name>:<your-local-registry-port>/redhat"
```

Additionally, ensure your certificate is present in the **additionalTrustBundle** field of the **install-config.yaml**.

Example install-config.yaml

```
additionalTrustBundle: |
-----BEGIN CERTIFICATE-----
ZZZZZZZZZZZZ
-----END CERTIFICATE-----
```

**IMPORTANT**

The **oc mirror** command creates a folder called **oc-mirror-workspace** with several outputs. This includes the **imageContentSourcePolicy.yaml** file that identifies all the mirrors you need for OpenShift Container Platform and your selected Operators.

- Generate the cluster manifests by running the following command:

```
$ openshift-install agent create cluster-manifests
```


This command updates the cluster manifests folder to include a **mirror** folder that contains your mirror configuration.

13.5.3. Preparing an Agent-based cluster deployment for the multicluster engine for Kubernetes Operator while connected

Create the required manifests for the multicluster engine Operator, the Local Storage Operator (LSO), and to deploy an agent-based OpenShift Container Platform cluster as a hub cluster.

Procedure

1. Create a sub-folder named **openshift** in the **<assets_directory>** folder. This sub-folder is used to store the extra manifests that will be applied during the installation to further customize the deployed cluster. The **<assets_directory>** folder contains all the assets including the **install-config.yaml** and **agent-config.yaml** files.



NOTE

The installer does not validate extra manifests.

2. For the multicluster engine, create the following manifests and save them in the **<assets_directory>/openshift** folder:

Example `mce_namespace.yaml`

```
apiVersion: v1
kind: Namespace
metadata:
  labels:
    openshift.io/cluster-monitoring: "true"
  name: multicluster-engine
```

Example `mce_operatorgroup.yaml`

```
apiVersion: operators.coreos.com/v1
kind: OperatorGroup
metadata:
  name: multicluster-engine-operatorgroup
  namespace: multicluster-engine
spec:
  targetNamespaces:
    - multicluster-engine
```

Example `mce_subscription.yaml`

```
apiVersion: operators.coreos.com/v1alpha1
kind: Subscription
metadata:
  name: multicluster-engine
  namespace: multicluster-engine
```

```
spec:
  channel: "stable-2.3"
  name: multicluster-engine
  source: redhat-operators
  sourceNamespace: openshift-marketplace
```



NOTE

You can install a distributed unit (DU) at scale with the Red Hat Advanced Cluster Management (RHACM) using the assisted installer (AI). These distributed units must be enabled in the hub cluster. The AI service requires persistent volumes (PVs), which are manually created.

- For the AI service, create the following manifests and save them in the **<assets_directory>/openshift** folder:

Example Iso_namespace.yaml


```
apiVersion: v1
kind: Namespace
metadata:
  annotations:
    openshift.io/cluster-monitoring: "true"
  name: openshift-local-storage
```

Example Iso_operatorgroup.yaml

```
apiVersion: operators.coreos.com/v1
kind: OperatorGroup
metadata:
  name: local-operator-group
  namespace: openshift-local-storage
spec:
  targetNamespaces:
    - openshift-local-storage
```

Example Iso_subscription.yaml

```
apiVersion: operators.coreos.com/v1alpha1
kind: Subscription
metadata:
  name: local-storage-operator
  namespace: openshift-local-storage
spec:
  installPlanApproval: Automatic
  name: local-storage-operator
  source: redhat-operators
  sourceNamespace: openshift-marketplace
```

**NOTE**

After creating all the manifests, your filesystem must display as follows:

Example Filesystem

```
<assets_directory>
├── install-config.yaml
├── agent-config.yaml
├── /openshift
│   ├── mce_namespace.yaml
│   ├── mce_operatorgroup.yaml
│   ├── mce_subscription.yaml
│   ├── lso_namespace.yaml
│   ├── lso_operatorgroup.yaml
│   └── lso_subscription.yaml
```

4. Create the agent ISO image by running the following command:

```
$ openshift-install agent create image --dir <assets_directory>
```

5. When the image is ready, boot the target machine and wait for the installation to complete.
6. To monitor the installation, run the following command:

```
$ openshift-install agent wait-for install-complete --dir <assets_directory>
```

**NOTE**


To configure a fully functional hub cluster, you must create the following manifests and manually apply them by running the command **\$ oc apply -f <manifest-name>**. The order of the manifest creation is important and where required, the waiting condition is displayed.

7. For the PVs that are required by the AI service, create the following manifests:

```
apiVersion: local.storage.openshift.io/v1
kind: LocalVolume
metadata:
  name: assisted-service
  namespace: openshift-local-storage
spec:
  logLevel: Normal
  managementState: Managed
  storageClassDevices:
    - devicePaths:
      - /dev/vda
      - /dev/vdb
    storageClassName: assisted-service
    volumeMode: Filesystem
```

8. Use the following command to wait for the availability of the PVs, before applying the subsequent manifests:

```
$ oc wait localvolume -n openshift-local-storage assisted-service --for condition=Available --
timeout 10m
```

**NOTE**

The `devicePath`` is an example and may vary depending on the actual hardware configuration used.

9. Create a manifest for a multicluster engine instance.

Example MultiClusterEngine.yaml

```
apiVersion: multicluster.openshift.io/v1
kind: MultiClusterEngine
metadata:
  name: multiclusterengine
spec: {}
```

10. Create a manifest to enable the AI service.

Example agentserviceconfig.yaml

```
apiVersion: agent-install.openshift.io/v1beta1
kind: AgentServiceConfig
metadata:
  name: agent
  namespace: assisted-installer
spec:
  databaseStorage:
    storageClassName: assisted-service
    accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 10Gi
  filesystemStorage:
    storageClassName: assisted-service
    accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 10Gi
```

11. Create a manifest to deploy subsequently spoke clusters.

Example clusterimageset.yaml

```
apiVersion: hive.openshift.io/v1
kind: ClusterImageSet
metadata:
```

```

name: "4.16"
spec:
  releaseImage: quay.io/openshift-release-dev/ocp-release:4.16.0-x86_64

```

12. Create a manifest to import the agent installed cluster (that hosts the multicluster engine and the Assisted Service) as the hub cluster.

Example `autoimport.yaml`

```

apiVersion: cluster.open-cluster-management.io/v1
kind: ManagedCluster
metadata:
  labels:
    local-cluster: "true"
    cloud: auto-detect
    vendor: auto-detect
  name: local-cluster
spec:
  hubAcceptsClient: true

```

13. Wait for the managed cluster to be created.

```

$ oc wait -n multicluster-engine managedclusters local-cluster --for
condition=ManagedClusterJoined=True --timeout 10m

```

Verification

- To confirm that the managed cluster installation is successful, run the following command:

```

$ oc get managedcluster
NAME          HUB ACCEPTED  MANAGED CLUSTER URLS      JOINED  AVAILABLE
AGE
local-cluster true          https://<your cluster url>:6443  True   True    77m

```

Additional resources

- [The Local Storage Operator](#)

13.6. INSTALLATION CONFIGURATION PARAMETERS FOR THE AGENT-BASED INSTALLER

Before you deploy an OpenShift Container Platform cluster using the Agent-based Installer, you provide parameters to customize your cluster and the platform that hosts it. When you create the **install-config.yaml** and **agent-config.yaml** files, you must provide values for the required parameters, and you can use the optional parameters to customize your cluster further.

13.6.1. Available installation configuration parameters

The following tables specify the required and optional installation configuration parameters that you can set as part of the Agent-based installation process.

These values are specified in the **install-config.yaml** file.

**NOTE**

These settings are used for installation only, and cannot be modified after installation.

13.6.1.1. Required configuration parameters

Required installation configuration parameters are described in the following table:

Table 13.7. Required parameters

Parameter	Description	Values
apiVersion:	The API version for the install-config.yaml content. The current version is v1 . The installation program may also support older API versions.	String
baseDomain:	The base domain of your cloud provider. The base domain is used to create routes to your OpenShift Container Platform cluster components. The full DNS name for your cluster is a combination of the baseDomain and metadata.name parameter values that uses the <metadata.name> . <baseDomain> format.	A fully-qualified domain or subdomain name, such as example.com .
metadata:	Kubernetes resource ObjectMeta , from which only the name parameter is consumed.	Object
metadata: name:	The name of the cluster. DNS records for the cluster are all subdomains of {{.metadata.name}} . {{.baseDomain}} . When you do not provide metadata.name through either the install-config.yaml or agent-config.yaml files, for example when you use only ZTP manifests, the cluster name is set to agent-cluster .	String of lowercase letters, hyphens (-), and periods (.), such as dev .

Parameter	Description	Values
<code>platform:</code>	The configuration for the specific platform upon which to perform the installation: baremetal , external , none , or vsphere .	Object
<code>pullSecret:</code>	Get a pull secret from Red Hat OpenShift Cluster Manager to authenticate downloading container images for OpenShift Container Platform components from services such as Quay.io.	<pre>{ "auths":{ "cloud.openshift.com":{ "auth":"b3Blb=", "email":"you@example.com" }, "quay.io":{ "auth":"b3Blb=", "email":"you@example.com" } } }</pre>

13.6.1.2. Network configuration parameters

You can customize your installation configuration based on the requirements of your existing network infrastructure. For example, you can expand the IP address block for the cluster network or provide different IP address blocks than the defaults.

- If you use the Red Hat OpenShift Networking OVN-Kubernetes network plugin, both IPv4 and IPv6 address families are supported.

If you configure your cluster to use both IP address families, review the following requirements:


- Both IP families must use the same network interface for the default gateway.
- Both IP families must have the default gateway.
- You must specify IPv4 and IPv6 addresses in the same order for all network configuration parameters. For example, in the following configuration IPv4 addresses are listed before IPv6 addresses.


```
networking:
  clusterNetwork:
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  - cidr: fd00:10:128::/56
    hostPrefix: 64
  serviceNetwork:
  - 172.30.0.0/16
  - fd00:172:16::/112
```

**NOTE**

Globalnet is not supported with Red Hat OpenShift Data Foundation disaster recovery solutions. For regional disaster recovery scenarios, ensure that you use a nonoverlapping range of private IP addresses for the cluster and service networks in each cluster.

Table 13.8. Network parameters

Parameter	Description	Values
<code>networking:</code>	The configuration for the cluster network.	Object  NOTE You cannot modify parameters specified by the networking object after installation.
<code>networking: networkType:</code>	The Red Hat OpenShift Networking network plugin to install.	OVNKubernetes. OVNKubernetes is a CNI plugin for Linux networks and hybrid networks that contain both Linux and Windows servers. The default value is OVNKubernetes .
<code>networking: clusterNetwork:</code>	The IP address blocks for pods. The default value is 10.128.0.0/14 with a host prefix of /23 . If you specify multiple IP address blocks, the blocks must not overlap.	An array of objects. For example: <code>networking: clusterNetwork: - cidr: 10.128.0.0/14 hostPrefix: 23 - cidr: fd01::/48 hostPrefix: 64</code>
<code>networking: clusterNetwork: cidr:</code>	Required if you use networking.clusterNetwork . An IP address block. If you use the OVN-Kubernetes network plugin, you can specify IPv4 and IPv6 networks.	An IP address block in Classless Inter-Domain Routing (CIDR) notation. The prefix length for an IPv4 block is between 0 and 32 . The prefix length for an IPv6 block is between 0 and 128 . For example, 10.128.0.0/14 or fd01::/48 .
<code>networking: clusterNetwork: hostPrefix:</code>	The subnet prefix length to assign to each individual node. For example, if hostPrefix is set to 23 then each node is assigned a /23 subnet out of the given cidr . A hostPrefix value of 23 provides 510 ($2^{(32 - 23)} - 2$) pod IP addresses.	A subnet prefix. For an IPv4 network the default value is 23 . For an IPv6 network the default value is 64 . The default value is also the minimum value for IPv6.

Parameter	Description	Values
<code>networking: serviceNetwork:</code>	<p>The IP address block for services. The default value is 172.30.0.0/16.</p> <p>The OVN-Kubernetes network plugins supports only a single IP address block for the service network.</p> <p>If you use the OVN-Kubernetes network plugin, you can specify an IP address block for both of the IPv4 and IPv6 address families.</p>	<p>An array with an IP address block in CIDR format. For example:</p> <pre>networking: serviceNetwork: - 172.30.0.0/16 - fd02::/112</pre>
<code>networking: machineNetwork:</code>	<p>The IP address blocks for machines.</p> <p>If you specify multiple IP address blocks, the blocks must not overlap.</p>	<p>An array of objects. For example:</p> <pre>networking: machineNetwork: - cidr: 10.0.0.0/16</pre>
<code>networking: machineNetwork: cidr:</code>	<p>Required if you use networking.machineNetwork. An IP address block. The default value is 10.0.0.0/16 for all platforms other than libvirt and IBM Power® Virtual Server. For libvirt, the default value is 192.168.126.0/24. For IBM Power® Virtual Server, the default value is 192.168.0.0/24.</p>	<p>An IP network block in CIDR notation.</p> <p>For example, 10.0.0.0/16 or fd00::/48.</p> <div style="display: flex; align-items: flex-start;">  <div> <p>NOTE</p> <p>Set the networking.machineNetwork to match the CIDR that the preferred NIC resides in.</p> </div> </div>


13.6.1.3. Optional configuration parameters


Optional installation configuration parameters are described in the following table:

Table 13.9. Optional parameters


Parameter	Description	Values
<code>additionalTrustBundle:</code>	<p>A PEM-encoded X.509 certificate bundle that is added to the nodes' trusted certificate store. This trust bundle may also be used when a proxy has been configured.</p>	String


Parameter	Description	Values
capabilities:	Controls the installation of optional core cluster components. You can reduce the footprint of your OpenShift Container Platform cluster by disabling optional components. For more information, see the "Cluster capabilities" page in <i>Installing</i> .	String array
capabilities: baselineCapabilitySet:	Selects an initial set of optional capabilities to enable. Valid values are None , v4.11 , v4.12 and vCurrent . The default value is vCurrent .	String
capabilities: additionalEnabledCapabilities:	Extends the set of optional capabilities beyond what you specify in baselineCapabilitySet . You may specify multiple capabilities in this parameter.	String array
cpuPartitioningMode:	Enables workload partitioning, which isolates OpenShift Container Platform services, cluster management workloads, and infrastructure pods to run on a reserved set of CPUs. Workload partitioning can only be enabled during installation and cannot be disabled after installation. While this field enables workload partitioning, it does not configure workloads to use specific CPUs. For more information, see the <i>Workload partitioning</i> page in the <i>Scalability and Performance</i> section.	None or AllNodes . None is the default value.
compute:	The configuration for the machines that comprise the compute nodes.	Array of MachinePool objects.
compute: architecture:	Determines the instruction set architecture of the machines in the pool. Currently, clusters with varied architectures are not supported. All pools must specify the same architecture. Valid values are amd64 , arm64 , ppc64le , and s390x .	String

Parameter	Description	Values
<code>compute: hyperthreading:</code>	<p>Whether to enable or disable simultaneous multithreading, or hyperthreading, on compute machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.</p>  <p>IMPORTANT</p> <p>If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.</p>	Enabled or Disabled
<code>compute: name:</code>	Required if you use compute . The name of the machine pool.	worker
<code>compute: platform:</code>	Required if you use compute . Use this parameter to specify the cloud provider to host the worker machines. This parameter value must match the controlPlane.platform parameter value.	baremetal , vsphere , or {}
<code>compute: replicas:</code>	The number of compute machines, which are also known as worker machines, to provision.	A positive integer greater than or equal to 2 . The default value is 3 .
<code>featureSet:</code>	Enables the cluster for a feature set. A feature set is a collection of OpenShift Container Platform features that are not enabled by default. For more information about enabling a feature set during installation, see "Enabling features using feature gates".	String. The name of the feature set to enable, such as TechPreviewNoUpgrade .
<code>controlPlane:</code>	The configuration for the machines that comprise the control plane.	Array of MachinePool objects.

Parameter	Description	Values
<code>controlPlane: architecture:</code>	Determines the instruction set architecture of the machines in the pool. Currently, clusters with varied architectures are not supported. All pools must specify the same architecture. Valid values are amd64 , arm64 , ppc64le , and s390x .	String
<code>controlPlane: hyperthreading:</code>	Whether to enable or disable simultaneous multithreading, or hyperthreading , on control plane machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.  IMPORTANT If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.	Enabled or Disabled
<code>controlPlane: name:</code>	Required if you use controlPlane . The name of the machine pool.	master
<code>controlPlane: platform:</code>	Required if you use controlPlane . Use this parameter to specify the cloud provider that hosts the control plane machines. This parameter value must match the compute.platform parameter value.	baremetal , vsphere , or {}
<code>controlPlane: replicas:</code>	The number of control plane machines to provision.	Supported values are 3 , or 1 when deploying single-node OpenShift.

Parameter	Description	Values
credentialsMode:	<p>The Cloud Credential Operator (CCO) mode. If no mode is specified, the CCO dynamically tries to determine the capabilities of the provided credentials, with a preference for mint mode on the platforms where multiple modes are supported.</p>	Mint, Passthrough, Manual or an empty string (""). ^[1]
fips:	<p>Enable or disable FIPS mode. The default is false (disabled). If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.</p> <div data-bbox="486 887 592 1935" style="background-color: black; color: white; padding: 10px; margin: 10px 0;"> <p>IMPORTANT</p> <p>To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see Installing the system in FIPS mode.</p> <p>When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.</p> </div>	false or true

Parameter	Description	NOTE	Values
		<p>If you are using Azure File storage, you cannot enable FIPS mode.</p>	
imageContentSources:	<p>Sources and repositories for the release-image content.</p>		<p>Array of objects. Includes a source and, optionally, mirrors, as described in the following rows of this table.</p>
imageContentSources: source:	<p>Required if you use imageContentSources. Specify the repository that users refer to, for example, in image pull specifications.</p>		<p>String</p>
imageContentSources: mirrors:	<p>Specify one or more repositories that may also contain the same images.</p>		<p>Array of strings</p>
publish:	<p>How to publish or expose the user-facing endpoints of your cluster, such as the Kubernetes API, OpenShift routes.</p>	<p>Internal or External. The default value is External.</p> <p>Setting this field to Internal is not supported on non-cloud platforms.</p> <div data-bbox="979 1442 1083 1697" style="background-color: black; width: 65px; height: 114px; margin: 10px 0;"></div> <p>IMPORTANT</p> <p>If the value of the field is set to Internal, the cluster will become non-functional. For more information, refer to BZ#1953035.</p>	

Parameter	Description	Values
<code>sshKey:</code>	<p>The SSH key to authenticate access to your cluster machines.</p>  <p>NOTE</p> <p>For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your ssh-agent process uses.</p>	For example, sshKey: ssh-ed25519 AAAA...

1. Not all CCO modes are supported for all cloud providers. For more information about CCO modes, see the "Managing cloud provider credentials" entry in the *Authentication and authorization* content.

13.6.1.4. Additional bare metal configuration parameters for the Agent-based Installer

Additional bare metal installation configuration parameters for the Agent-based Installer are described in the following table:



NOTE

These fields are not used during the initial provisioning of the cluster, but they are available to use once the cluster has been installed. Configuring these fields at install time eliminates the need to set them as a Day 2 operation.

Table 13.10. Additional bare metal parameters

Parameter	Description	Values
<code>platform: baremetal: clusterProvisioningI P:</code>	<p>The IP address within the cluster where the provisioning services run. Defaults to the third IP address of the provisioning subnet. For example, 172.22.0.3 or 2620:52:0:1307::3.</p>	IPv4 or IPv6 address.

Parameter	Description	Values
<p>platform: baremetal:</p> <p>provisioningNetwork:</p>	<p>The provisioningNetwork configuration setting determines whether the cluster uses the provisioning network. If it does, the configuration setting also determines if the cluster manages the network.</p> <p>Managed: Default. Set this parameter to Managed to fully manage the provisioning network, including DHCP, TFTP, and so on.</p> <p>Disabled: Set this parameter to Disabled to disable the requirement for a provisioning network. When set to Disabled, you can use only virtual media based provisioning on Day 2. If Disabled and using power management, BMCs must be accessible from the bare-metal network. If Disabled, you must provide two IP addresses on the bare-metal network that are used for the provisioning services.</p>	<p>Managed or Disabled.</p>
<p>platform: baremetal:</p> <p>provisioningMACAddress:</p>	<p>The MAC address within the cluster where provisioning services run.</p>	<p>MAC address.</p>
<p>platform: baremetal:</p> <p>provisioningNetworkCIDR:</p>	<p>The CIDR for the network to use for provisioning. This option is required when not using the default address range on the provisioning network.</p>	<p>Valid CIDR, for example 10.0.0.0/16.</p>
<p>platform: baremetal:</p> <p>provisioningNetworkInterface:</p>	<p>The name of the network interface on nodes connected to the provisioning network. Use the bootMACAddress configuration setting to enable Ironic to identify the IP address of the NIC instead of using the provisioningNetworkInterface configuration setting to identify the name of the NIC.</p>	<p>String.</p>

Parameter	Description	Values
<pre>platform: baremetal: provisioningDHCP Range:</pre>	<p>Defines the IP range for nodes on the provisioning network, for example 172.22.0.10,172.22.0.254.</p>	IP address range.
<pre>platform: baremetal: hosts:</pre>	Configuration for bare metal hosts.	Array of host configuration objects.
<pre>platform: baremetal: hosts: name:</pre>	The name of the host.	String.
<pre>platform: baremetal: hosts: bootMACAddress:</pre>	The MAC address of the NIC used for provisioning the host.	MAC address.
<pre>platform: baremetal: hosts: bmc:</pre>	Configuration for the host to connect to the baseboard management controller (BMC).	Dictionary of BMC configuration objects.
<pre>platform: baremetal: hosts: bmc: username:</pre>	The username for the BMC.	String.
<pre>platform: baremetal: hosts: bmc: password:</pre>	Password for the BMC.	String.

Parameter	Description	Values
<ul style="list-style-type: none"> platform: baremetal: hosts: bmc: address: 	<p>The URL for communicating with the host's BMC controller. The address configuration setting specifies the protocol. For example, redfish+http://10.10.10.1:8000/redfish/v1/Systems/1234 enables Redfish. For more information, see "BMC addressing" in the "Deploying installer-provisioned clusters on bare metal" section.</p>	URL.
<ul style="list-style-type: none"> platform: baremetal: hosts: bmc: <p>disableCertificateVerification:</p>	<p>redfish and redfish-virtualmedia need this parameter to manage BMC addresses. The value should be True when using a self-signed certificate for BMC addresses.</p>	Boolean.


13.6.1.5. Additional VMware vSphere configuration parameters

Additional VMware vSphere configuration parameters are described in the following table:

Table 13.11. Additional VMware vSphere cluster parameters

Parameter	Description	Values
<ul style="list-style-type: none"> platform: vsphere: 	<p>Describes your account on the cloud platform that hosts your cluster. You can use the parameter to customize the platform. If you provide additional configuration settings for compute and control plane machines in the machine pool, the parameter is not required. You can only specify one vCenter server for your OpenShift Container Platform cluster.</p>	A dictionary of vSphere configuration objects
<ul style="list-style-type: none"> platform: vsphere: failureDomains: 	<p>Establishes the relationships between a region and zone. You define a failure domain by using vCenter objects, such as a datastore object. A failure domain defines the vCenter location for OpenShift Container Platform cluster nodes.</p>	An array of failure domain configuration objects.
<ul style="list-style-type: none"> platform: vsphere: failureDomains: name: 	<p>The name of the failure domain.</p>	String

Parameter	Description	Values
<pre>platform: vsphere: failureDomains: region:</pre>	<p>If you define multiple failure domains for your cluster, you must attach the tag to each vCenter datacenter. To define a region, use a tag from the openshift-region tag category. For a single vSphere datacenter environment, you do not need to attach a tag, but you must enter an alphanumeric value, such as datacenter, for the parameter.</p>	String
<pre>platform: vsphere: failureDomains: server:</pre>	<p>Specifies the fully-qualified hostname or IP address of the VMware vCenter server, so that a client can access failure domain resources. You must apply the server role to the vSphere vCenter server location.</p>	String
<pre>platform: vsphere: failureDomains: zone:</pre>	<p>If you define multiple failure domains for your cluster, you must attach a tag to each vCenter cluster. To define a zone, use a tag from the openshift-zone tag category. For a single vSphere datacenter environment, you do not need to attach a tag, but you must enter an alphanumeric value, such as cluster, for the parameter.</p>	String
<pre>platform: vsphere: failureDomains: topology: computeCluster:</pre>	<p>The path to the vSphere compute cluster.</p>	String
<pre>platform: vsphere: failureDomains: topology: datacenter:</pre>	<p>Lists and defines the datacenters where OpenShift Container Platform virtual machines (VMs) operate. The list of datacenters must match the list of datacenters specified in the vcenters field.</p>	String

Parameter	Description	Values
<pre>platform: vsphere: failureDomains: topology: datastore:</pre>	<p>The path to the vSphere datastore that holds virtual machine files, templates, and ISO images.</p>  <p>IMPORTANT</p> <p>You can specify the path of any datastore that exists in a datastore cluster. By default, Storage vMotion is automatically enabled for a datastore cluster. Red Hat does not support Storage vMotion, so you must disable Storage vMotion to avoid data loss issues for your OpenShift Container Platform cluster.</p> <p>If you must specify VMs across multiple datastores, use a datastore object to specify a failure domain in your cluster's install-config.yaml configuration file. For more information, see "VMware vSphere region and zone enablement".</p>	String
<pre>platform: vsphere: failureDomains: topology: folder:</pre>	<p>Optional: The absolute path of an existing folder where the user creates the virtual machines, for example, /<datacenter_name>/vm/<folder_name>/<sub folder_name>.</p>	String
<pre>platform: vsphere: failureDomains: topology: networks:</pre>	<p>Lists any network in the vCenter instance that contains the virtual IP addresses and DNS records that you configured.</p>	String

Parameter	Description	Values
<pre>platform: vsphere: failureDomains: topology: resourcePool:</pre>	Optional: The absolute path of an existing resource pool where the installation program creates the virtual machines, for example, <code>/<datacenter_name>/host/<cluster_name>/Resources/<resource_pool_name>/<optional_nested_resource_pool_name></code> .	String
<pre>platform: vsphere: failureDomains: topology template:</pre>	Specifies the absolute path to a pre-existing Red Hat Enterprise Linux CoreOS (RHCOS) image template or virtual machine. The installation program can use the image template or virtual machine to quickly install RHCOS on vSphere hosts. Consider using this parameter as an alternative to uploading an RHCOS image on vSphere hosts. This parameter is available for use only on installer-provisioned infrastructure.	String
<pre>platform: vsphere: vcenters:</pre>	Configures the connection details so that services can communicate with a vCenter server. Currently, only a single vCenter server is supported.	An array of vCenter configuration objects.
<pre>platform: vsphere: vcenters: datacenters:</pre>	Lists and defines the datacenters where OpenShift Container Platform virtual machines (VMs) operate. The list of datacenters must match the list of datacenters specified in the failureDomains field.	String
<pre>platform: vsphere: vcenters: password:</pre>	The password associated with the vSphere user.	String
<pre>platform: vsphere: vcenters: port:</pre>	The port number used to communicate with the vCenter server.	Integer
<pre>platform: vsphere: vcenters: server:</pre>	The fully qualified host name (FQHN) or IP address of the vCenter server.	String

Parameter	Description	Values
platform: vsphere: vcenters: user:	The username associated with the vSphere user.	String

13.6.1.6. Deprecated VMware vSphere configuration parameters

In OpenShift Container Platform 4.13, the following vSphere configuration parameters are deprecated. You can continue to use these parameters, but the installation program does not automatically specify these parameters in the **install-config.yaml** file.

The following table lists each deprecated vSphere configuration parameter:

Table 13.12. Deprecated VMware vSphere cluster parameters

Parameter	Description	Values
platform: vsphere: cluster:	The vCenter cluster to install the OpenShift Container Platform cluster in.	String
platform: vsphere: datacenter:	Defines the datacenter where OpenShift Container Platform virtual machines (VMs) operate.	String
platform: vsphere: defaultDatastore:	The name of the default datastore to use for provisioning volumes.	String
platform: vsphere: folder:	Optional: The absolute path of an existing folder where the installation program creates the virtual machines. If you do not provide this value, the installation program creates a folder that is named with the infrastructure ID in the data center virtual machine folder.	String, for example, /<datacenter_name>/vm/<folder_name>/<subfolder_name> .
platform: vsphere: password:	The password for the vCenter user name.	String

Parameter	Description	Values
platform: vsphere: resourcePool:	Optional: The absolute path of an existing resource pool where the installation program creates the virtual machines. If you do not specify a value, the installation program installs the resources in the root of the cluster under <code>/<datacenter_name>/host/<cluster_name>/Resources</code> .	String, for example, <code>/<datacenter_name>/host/<cluster_name>/Resources/<resource_pool_name>/<optional_nested_resource_pool_name></code> .
platform: vsphere: username:	The user name to use to connect to the vCenter instance with. This user must have at least the roles and privileges that are required for static or dynamic persistent volume provisioning in vSphere.	String
platform: vsphere: vCenter:	The fully-qualified hostname or IP address of a vCenter server.	String

Additional resources

- [BMC addressing](#)
- [Configuring regions and zones for a VMware vCenter](#)

13.6.2. Available Agent configuration parameters

The following tables specify the required and optional Agent configuration parameters that you can set as part of the Agent-based installation process.

These values are specified in the `agent-config.yaml` file.



NOTE

These settings are used for installation only, and cannot be modified after installation.

13.6.2.1. Required configuration parameters

Required Agent configuration parameters are described in the following table:

Table 13.13. Required parameters

Parameter	Description	Values
-----------	-------------	--------

Parameter	Description	Values
<code>apiVersion:</code>	The API version for the agent-config.yaml content. The current version is v1beta1 . The installation program might also support older API versions.	String
<code>metadata:</code>	Kubernetes resource ObjectMeta , from which only the name parameter is consumed.	Object
<code>metadata: name:</code>	The name of the cluster. DNS records for the cluster are all subdomains of {{.metadata.name}} . {{.baseDomain}} . The value entered in the agent-config.yaml file is ignored, and instead the value specified in the install-config.yaml file is used. When you do not provide metadata.name through either the install-config.yaml or agent-config.yaml files, for example when you use only ZTP manifests, the cluster name is set to agent-cluster .	String of lowercase letters and hyphens (-), such as dev .

13.6.2.2. Optional configuration parameters

Optional Agent configuration parameters are described in the following table:

Table 13.14. Optional parameters

Parameter	Description	Values
-----------	-------------	--------

Parameter	Description	Values
<code>rendezvousIP:</code>	The IP address of the node that performs the bootstrapping process as well as running the assisted-service component. You must provide the rendezvous IP address when you do not specify at least one host's IP address in the networkConfig parameter. If this address is not provided, one IP address is selected from the provided hosts' networkConfig .	IPv4 or IPv6 address.
<code>bootArtifactsBaseURL:</code>	The URL of the server to upload Preboot Execution Environment (PXE) assets to when using the Agent-based Installer to generate an iPXE script. For more information, see "Preparing PXE assets for OpenShift Container Platform".	String.
<code>additionalNTPSources:</code>	A list of Network Time Protocol (NTP) sources to be added to all cluster hosts, which are added to any NTP sources that are configured through other means.	List of hostnames or IP addresses.
<code>hosts:</code>	Host configuration. An optional list of hosts. The number of hosts defined must not exceed the total number of hosts defined in the install-config.yaml file, which is the sum of the values of the compute.replicas and controlPlane.replicas parameters.	An array of host configuration objects.
<code>hosts: hostname:</code>	Hostname. Overrides the hostname obtained from either the Dynamic Host Configuration Protocol (DHCP) or a reverse DNS lookup. Each host must have a unique hostname supplied by one of these methods, although configuring a hostname through this parameter is optional.	String.
<code>hosts: interfaces:</code>	Provides a table of the name and MAC address mappings for the interfaces on the host. If a NetworkConfig section is provided in the agent-config.yaml file, this table must be included and the values must match the mappings provided in the NetworkConfig section.	An array of host configuration objects.

Parameter	Description	Values
hosts: interfaces: name:	The name of an interface on the host.	String.
hosts: interfaces: macAddress:	The MAC address of an interface on the host.	A MAC address such as the following example: 00-B0-D0-63-C2-26 .
hosts: role:	Defines whether the host is a master or worker node. If no role is defined in the agent-config.yaml file, roles will be assigned at random during cluster installation.	master or worker .
hosts: rootDeviceHints:	Enables provisioning of the Red Hat Enterprise Linux CoreOS (RHCOS) image to a particular device. The installation program examines the devices in the order it discovers them, and compares the discovered values with the hint values. It uses the first discovered device that matches the hint value. This is the device that the operating system is written on during installation.	A dictionary of key-value pairs. For more information, see "Root device hints" in the "Setting up the environment for an OpenShift installation" page.
hosts: rootDeviceHints: deviceName:	The name of the device the RHCOS image is provisioned to.	String.
hosts: networkConfig:	The host network definition. The configuration must match the Host Network Management API defined in the nmstate documentation .	A dictionary of host network configuration objects.

Additional resources

- [Preparing PXE assets for OpenShift Container Platform](#)
- [Root device hints](#)

CHAPTER 14. INSTALLING ON A SINGLE NODE

14.1. PREPARING TO INSTALL ON A SINGLE NODE

14.1.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You have read the documentation on [selecting a cluster installation method and preparing it for users](#).

14.1.2. About OpenShift on a single node

You can create a single-node cluster with standard installation methods. OpenShift Container Platform on a single node is a specialized installation that requires the creation of a special Ignition configuration file. The primary use case is for edge computing workloads, including intermittent connectivity, portable clouds, and 5G radio access networks (RAN) close to a base station. The major tradeoff with an installation on a single node is the lack of high availability.



IMPORTANT

The use of OpenShiftSDN with single-node OpenShift is not supported. OVN-Kubernetes is the default network plugin for single-node OpenShift deployments.

14.1.3. Requirements for installing OpenShift on a single node


Installing OpenShift Container Platform on a single node alleviates some of the requirements for high availability and large scale clusters. However, you must address the following requirements:

- **Administration host:** You must have a computer to prepare the ISO, to create the USB boot drive, and to monitor the installation.



NOTE

For the **ppc64le** platform, the host should prepare the ISO, but does not need to create the USB boot drive. The ISO can be mounted to PowerVM directly.




NOTE

ISO is not required for IBM Z[®] installations.

- **CPU Architecture:** Installing OpenShift Container Platform on a single node supports **x86_64**, **arm64**, **ppc64le**, and **s390x** CPU architectures.
- **Supported platforms:** Installing OpenShift Container Platform on a single node is supported on bare metal, vSphere, AWS cloud, Red Hat OpenStack, OpenShift Virtualization, IBM Power[®], and IBM Z[®] platforms.
- **Production-grade server:** Installing OpenShift Container Platform on a single node requires a server with sufficient resources to run OpenShift Container Platform services and a production workload.

Table 14.1. Minimum resource requirements

Profile	vCPU	Memory	Storage
Minimum	8 vCPU cores	16 GB of RAM	120 GB

**NOTE**

- One vCPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio:
(threads per core × cores) × sockets = vCPUs
- Adding Operators during the installation process might increase the minimum resource requirements.

The server must have a Baseboard Management Controller (BMC) when booting with virtual media.


**NOTE**

BMC is not supported on IBM Z® and IBM Power®.

- Networking:** The server must have access to the internet or access to a local registry if it is not connected to a routable network. The server must have a DHCP reservation or a static IP address for the Kubernetes API, ingress route, and cluster node domain names. You must configure the DNS to resolve the IP address to each of the following fully qualified domain names (FQDN):

Table 14.2. Required DNS records

Usage	FQDN	Description
Kubernetes API	api.<cluster_name>.<base_domain>	Add a DNS A/AAAA or CNAME record. This record must be resolvable by both clients external to the cluster and within the cluster.
Internal API	api-int.<cluster_name>.<base_domain>	Add a DNS A/AAAA or CNAME record when creating the ISO manually. This record must be resolvable by nodes within the cluster.
Ingress route	*.apps.<cluster_name>.<base_domain>	Add a wildcard DNS A/AAAA or CNAME record that targets the node. This record must be resolvable by both clients external to the cluster and within the cluster.

**IMPORTANT**

Without persistent IP addresses, communications between the **apiserver** and **etcd** might fail.

14.2. INSTALLING OPENSIFT ON A SINGLE NODE

You can install single-node OpenShift using the web-based Assisted Installer and a discovery ISO that you generate using the Assisted Installer. You can also install single-node OpenShift by using **coreos-installer** to generate the installation ISO.

14.2.1. Installing single-node OpenShift using the Assisted Installer

To install OpenShift Container Platform on a single node, use the web-based Assisted Installer wizard to guide you through the process and manage the installation.

See the [Assisted Installer for OpenShift Container Platform](#) documentation for details and configuration options.

14.2.1.1. Generating the discovery ISO with the Assisted Installer

Installing OpenShift Container Platform on a single node requires a discovery ISO, which the Assisted Installer can generate.

Procedure

1. On the administration host, open a browser and navigate to [Red Hat OpenShift Cluster Manager](#).
2. Click **Create Cluster** to create a new cluster.
3. In the **Cluster name** field, enter a name for the cluster.
4. In the **Base domain** field, enter a base domain. For example:

```
example.com
```

All DNS records must be subdomains of this base domain and include the cluster name, for example:

```
<cluster-name>.example.com
```

**NOTE**

You cannot change the base domain or cluster name after cluster installation.

5. Select **Install single node OpenShift (SNO)** and complete the rest of the wizard steps. Download the discovery ISO.
6. Make a note of the discovery ISO URL for installing with virtual media.

**NOTE**

If you enable OpenShift Virtualization during this process, you must have a second local storage device of at least 50GiB for your virtual machines.

Additional resources

- [Persistent storage using logical volume manager storage](#)
- [What you can do with OpenShift Virtualization](#)

14.2.1.2. Installing single-node OpenShift with the Assisted Installer

Use the Assisted Installer to install the single-node cluster.

Procedure

1. Attach the RHCOS discovery ISO to the target host.
2. Configure the boot drive order in the server BIOS settings to boot from the attached discovery ISO and then reboot the server.
3. On the administration host, return to the browser. Wait for the host to appear in the list of discovered hosts. If necessary, reload the [Assisted Clusters](#) page and select the cluster name.
4. Complete the install wizard steps. Add networking details, including a subnet from the available subnets. Add the SSH public key if necessary.
5. Monitor the installation's progress. Watch the cluster events. After the installation process finishes writing the operating system image to the server's hard disk, the server restarts.
6. Remove the discovery ISO, and reset the server to boot from the installation drive. The server restarts several times automatically, deploying the control plane.

Additional resources

- [Creating a bootable ISO image on a USB drive](#)
- [Booting from an HTTP-hosted ISO image using the Redfish API](#)
- [Adding worker nodes to single-node OpenShift clusters](#)

14.2.2. Installing single-node OpenShift manually

To install OpenShift Container Platform on a single node, first generate the installation ISO, and then boot the server from the ISO. You can monitor the installation using the **openshift-install** installation program.

Additional resources

- [Networking requirements for user-provisioned infrastructure](#)
- [User-provisioned DNS requirements](#)
- [Configuring DHCP or static IP addresses](#)

14.2.2.1. Generating the installation ISO with coreos-installer

Installing OpenShift Container Platform on a single node requires an installation ISO, which you can generate with the following procedure.

Prerequisites

- Install **podman**.



NOTE

See "Requirements for installing OpenShift on a single node" for networking requirements, including DNS records.

Procedure

1. Set the OpenShift Container Platform version:

```
$ OCP_VERSION=<ocp_version> 1
```

- 1 Replace **<ocp_version>** with the current version, for example, **latest-4.16**

2. Set the host architecture:

```
$ ARCH=<architecture> 1
```

- 1 Replace **<architecture>** with the target host architecture, for example, **aarch64** or **x86_64**.

3. Download the OpenShift Container Platform client (**oc**) and make it available for use by entering the following commands:

```
$ curl -k https://mirror.openshift.com/pub/openshift-  
v4/clients/ocp/$OCP_VERSION/openshift-client-linux.tar.gz -o oc.tar.gz
```

```
$ tar zxf oc.tar.gz
```

```
$ chmod +x oc
```

4. Download the OpenShift Container Platform installer and make it available for use by entering the following commands:

```
$ curl -k https://mirror.openshift.com/pub/openshift-  
v4/clients/ocp/$OCP_VERSION/openshift-install-linux.tar.gz -o openshift-install-linux.tar.gz
```

```
$ tar zxvf openshift-install-linux.tar.gz
```

```
$ chmod +x openshift-install
```

5. Retrieve the RHCOS ISO URL by running the following command:

```
$ ISO_URL=$(./openshift-install coreos print-stream-json | grep location | grep $ARCH | grep iso | cut -d\" -f4)
```

6. Download the RHCOS ISO:

```
$ curl -L $ISO_URL -o rhcos-live.iso
```

7. Prepare the **install-config.yaml** file:

```
apiVersion: v1
baseDomain: <domain> 1
compute:
- name: worker
  replicas: 0 2
controlPlane:
  name: master
  replicas: 1 3
metadata:
  name: <name> 4
networking: 5
  clusterNetwork:
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  machineNetwork:
  - cidr: 10.0.0.0/16 6
  networkType: OVNKubernetes
  serviceNetwork:
  - 172.30.0.0/16
platform:
  none: {}
bootstrapInPlace:
  installationDisk: /dev/disk/by-id/<disk_id> 7
pullSecret: '<pull_secret>' 8
sshKey: |
  <ssh_key> 9
```

- 1 Add the cluster domain name.
- 2 Set the **compute** replicas to **0**. This makes the control plane node schedulable.
- 3 Set the **controlPlane** replicas to **1**. In conjunction with the previous **compute** setting, this setting ensures the cluster runs on a single node.
- 4 Set the **metadata** name to the cluster name.
- 5 Set the **networking** details. OVN-Kubernetes is the only allowed network plugin type for single-node clusters.
- 6 Set the **cidr** value to match the subnet of the single-node OpenShift cluster.
- 7 Set the path to the installation disk drive, for example, **/dev/disk/by-id/wwn-0x64cd98f04fde100024684cf3034da5c2**.
- 8 Copy the [pull secret from Red Hat OpenShift Cluster Manager](#) and add the contents to this configuration setting.

this configuration setting.

9. Add the public SSH key from the administration host so that you can log in to the cluster after installation.

8. Generate OpenShift Container Platform assets by running the following commands:

```
$ mkdir ocp
```

```
$ cp install-config.yaml ocp
```

```
$ ./openshift-install --dir=ocp create single-node-ignition-config
```

9. Embed the ignition data into the RHCOS ISO by running the following commands:

```
$ alias coreos-installer='podman run --privileged --pull always --rm \
-v /dev:/dev -v /run/udev:/run/udev -v $PWD:/data \
-w /data quay.io/coreos/coreos-installer:release'
```

```
$ coreos-installer iso ignition embed -fi ocp/bootstrap-in-place-for-live-iso.ign rhcos-live.iso
```

Additional resources

- See [Requirements for installing OpenShift on a single node](#) for more information about installing OpenShift Container Platform on a single node.
- See [Enabling cluster capabilities](#) for more information about enabling cluster capabilities that were disabled prior to installation.
- See [Optional cluster capabilities in OpenShift Container Platform 4.16](#) for more information about the features provided by each capability.

14.2.2.2. Monitoring the cluster installation using openshift-install

Use **openshift-install** to monitor the progress of the single-node cluster installation.

Procedure

1. Attach the modified RHCOS installation ISO to the target host.
2. Configure the boot drive order in the server BIOS settings to boot from the attached discovery ISO and then reboot the server.
3. On the administration host, monitor the installation by running the following command:

```
$ ./openshift-install --dir=ocp wait-for install-complete
```

The server restarts several times while deploying the control plane.

Verification

- After the installation is complete, check the environment by running the following command:
 -

```
$ export KUBECONFIG=ocp/auth/kubeconfig
```

```
$ oc get nodes
```

Example output

```
NAME                STATUS ROLES    AGE  VERSION
control-plane.example.com Ready  master,worker 10m  v1.29.4
```

Additional resources

- [Creating a bootable ISO image on a USB drive](#)
- [Booting from an HTTP-hosted ISO image using the Redfish API](#)
- [Adding worker nodes to single-node OpenShift clusters](#)

14.2.3. Installing single-node OpenShift on cloud providers

14.2.3.1. Additional requirements for installing single-node OpenShift on a cloud provider

The documentation for installer-provisioned installation on cloud providers is based on a high availability cluster consisting of three control plane nodes. When referring to the documentation, consider the differences between the requirements for a single-node OpenShift cluster and a high availability cluster.

- A high availability cluster requires a temporary bootstrap machine, three control plane machines, and at least two compute machines. For a single-node OpenShift cluster, you need only a temporary bootstrap machine and one cloud instance for the control plane node and no worker nodes.
- The minimum resource requirements for high availability cluster installation include a control plane node with 4 vCPUs and 100GB of storage. For a single-node OpenShift cluster, you must have a minimum of 8 vCPU cores and 120GB of storage.
- The **controlPlane.replicas** setting in the **install-config.yaml** file should be set to **1**.
- The **compute.replicas** setting in the **install-config.yaml** file should be set to **0**. This makes the control plane node schedulable.

14.2.3.2. Supported cloud providers for single-node OpenShift

The following table contains a list of supported cloud providers and CPU architectures.

Table 14.3. Supported cloud providers

Cloud provider	CPU architecture
Amazon Web Service (AWS)	x86_64 and AArch64
Microsoft Azure	x86_64
Google Cloud Platform (GCP)	x86_64 and AArch64

14.2.3.3. Installing single-node OpenShift on AWS

Installing a single-node cluster on AWS requires installer-provisioned installation using the "Installing a cluster on AWS with customizations" procedure.

Additional resources

- [Installing a cluster on AWS with customizations](#)

14.2.3.4. Installing single-node OpenShift on Azure

Installing a single node cluster on Azure requires installer-provisioned installation using the "Installing a cluster on Azure with customizations" procedure.

Additional resources

- [Installing a cluster on Azure with customizations](#)

14.2.3.5. Installing single-node OpenShift on GCP

Installing a single node cluster on GCP requires installer-provisioned installation using the "Installing a cluster on GCP with customizations" procedure.

Additional resources

- [Installing a cluster on GCP with customizations](#)

14.2.4. Creating a bootable ISO image on a USB drive

You can install software using a bootable USB drive that contains an ISO image. Booting the server with the USB drive prepares the server for the software installation.

Procedure

1. On the administration host, insert a USB drive into a USB port.
2. Create a bootable USB drive, for example:

```
# dd if=<path_to_iso> of=<path_to_usb> status=progress
```

where:

<path_to_iso>

is the relative path to the downloaded ISO file, for example, **rhcos-live.iso**.

<path_to_usb>

is the location of the connected USB drive, for example, **/dev/sdb**.

After the ISO is copied to the USB drive, you can use the USB drive to install software on the server.

14.2.5. Booting from an HTTP-hosted ISO image using the Redfish API

You can provision hosts in your network using ISOs that you install using the Redfish Baseboard Management Controller (BMC) API.



NOTE

This example procedure demonstrates the steps on a Dell server.



IMPORTANT

Ensure that you have the latest firmware version of iDRAC that is compatible with your hardware. If you have any issues with the hardware or firmware, you must contact the provider.

Prerequisites

- Download the installation Red Hat Enterprise Linux CoreOS (RHCOS) ISO.
- Use a Dell PowerEdge server that is compatible with iDRAC9.

Procedure

1. Copy the ISO file to an HTTP server accessible in your network.
2. Boot the host from the hosted ISO file, for example:
 - a. Call the Redfish API to set the hosted ISO as the **VirtualMedia** boot media by running the following command:

```
$ curl -k -u <bmc_username>:<bmc_password> -d '{"Image": "<hosted_iso_file>",
"Inserted": true}' -H "Content-Type: application/json" -X POST
<host_bmc_address>/redfish/v1/Managers/iDRAC.Embedded.1/VirtualMedia/CD/Actions/Vi
rtualMedia.InsertMedia
```

Where:

<bmc_username>:<bmc_password>

Is the username and password for the target host BMC.

<hosted_iso_file>

Is the URL for the hosted installation ISO, for example:

<http://webserver.example.com/rhcos-live-minimal.iso>. The ISO must be accessible from the target host machine.

<host_bmc_address>

Is the BMC IP address of the target host machine.

- b. Set the host to boot from the **VirtualMedia** device by running the following command:

```
$ curl -k -u <bmc_username>:<bmc_password> -X PATCH -H 'Content-Type:
application/json' -d '{"Boot": {"BootSourceOverrideTarget": "Cd",
"BootSourceOverrideMode": "UEFI", "BootSourceOverrideEnabled": "Once"}}'
<host_bmc_address>/redfish/v1/Systems/System.Embedded.1
```

- c. Reboot the host:

```
$ curl -k -u <bmc_username>:<bmc_password> -d '{"ResetType": "ForceRestart"}' -H
'Content-type: application/json' -X POST
<host_bmc_address>/redfish/v1/Systems/System.Embedded.1/Actions/ComputerSystem.Reset
```

- d. Optional: If the host is powered off, you can boot it using the **{"ResetType": "On"}** switch. Run the following command:

```
$ curl -k -u <bmc_username>:<bmc_password> -d '{"ResetType": "On"}' -H 'Content-
type: application/json' -X POST
<host_bmc_address>/redfish/v1/Systems/System.Embedded.1/Actions/ComputerSystem.Reset
```

14.2.6. Creating a custom live RHCOS ISO for remote server access

In some cases, you cannot attach an external disk drive to a server, however, you need to access the server remotely to provision a node. It is recommended to enable SSH access to the server. You can create a live RHCOS ISO with SSHd enabled and with predefined credentials so that you can access the server after it boots.

Prerequisites

- You installed the **butane** utility.

Procedure

- Download the **coreos-installer** binary from the **coreos-installer** image [mirror](#) page.
- Download the latest live RHCOS ISO from [mirror.openshift.com](#).
- Create the **embedded.yaml** file that the **butane** utility uses to create the Ignition file:

```
variant: openshift
version: 4.16.0
metadata:
  name: sshd
  labels:
    machineconfiguration.openshift.io/role: worker
passwd:
users:
  - name: core 1
    ssh_authorized_keys:
      - '<ssh_key>'
```

- 1** The **core** user has sudo privileges.

- Run the **butane** utility to create the Ignition file using the following command:

```
$ butane -pr embedded.yaml -o embedded.ign
```

- After the Ignition file is created, you can include the configuration in a new live RHCOS ISO, which is named **rhcos-sshd-4.16.0-x86_64-live.x86_64.iso**, with the **coreos-installer** utility:


```
$ coreos-installer iso ignition embed -i embedded.ign rhcos-4.16.0-x86_64-live.x86_64.iso -o
rhcos-sshd-4.16.0-x86_64-live.x86_64.iso
```

Verification

- Check that the custom live ISO can be used to boot the server by running the following command:

```
# coreos-installer iso ignition show rhcos-sshd-4.16.0-x86_64-live.x86_64.iso
```

Example output

```
{
  "ignition": {
    "version": "3.2.0"
  },
  "passwd": {
    "users": [
      {
        "name": "core",
        "sshAuthorizedKeys": [
          "ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAQACzNzG8AlzIDAhpyENpK2qKiTT8EbRWOrz7NXjRzo
pbPu215mocaJgjjwJjh1cYhgPhpAp6M/ttTk714OI7g4588ApX4bwJep6oWTU35LkY8ZxkGVPJL
8kVITdKQviDv3XX12I4QfnDom4tm4gVbRH0gNT1wzhnLP+LKYM2Ohr9D7p9NBnAdro6k++X
WgkDeijLRUTwdEyWunldW1f8G0Mg8Y1Xzr13BUo3+8aey7HLKJMDtobkz/C8ESYA/f7HJc5Fx
F0XbapWWovSSDJrr9OmIL9f4TfE+cQk3s+eoKiz2bgNPRgEEwihVbGsCN4grA+RzLCAOpec+
2dTJrQvFqsD alosadag@sonnelicht.local"
        ]
      }
    ]
  }
}
```

14.2.7. Installing single-node OpenShift with IBM Z and IBM LinuxONE

Installing a single-node cluster on IBM Z® and IBM® LinuxONE requires user-provisioned installation using one of the following procedures:

- [Installing a cluster with z/VM on IBM Z® and IBM® LinuxONE](#)
- [Installing a cluster with RHEL KVM on IBM Z® and IBM® LinuxONE](#)
- [Installing a cluster in an LPAR on IBM Z® and IBM® LinuxONE](#)




NOTE

Installing a single-node cluster on IBM Z® simplifies installation for development and test environments and requires less resource requirements at entry level.

Hardware requirements

- The equivalent of two Integrated Facilities for Linux (IFL), which are SMT2 enabled, for each cluster.
- At least one network connection to both connect to the **LoadBalancer** service and to serve data for traffic outside the cluster.



NOTE

You can use dedicated or shared IFLs to assign sufficient compute resources. Resource sharing is one of the key strengths of IBM Z®. However, you must adjust capacity correctly on each hypervisor layer and ensure sufficient resources for every OpenShift Container Platform cluster.

14.2.7.1. Installing single-node OpenShift with z/VM on IBM Z and IBM LinuxONE

Prerequisites

- You have installed **podman**.

Procedure

1. Set the OpenShift Container Platform version by running the following command:

```
$ OCP_VERSION=<ocp_version> 1
```

- 1 Replace **<ocp_version>** with the current version. For example, **latest-4.16**.

2. Set the host architecture by running the following command:

```
$ ARCH=<architecture> 1
```

- 1 Replace **<architecture>** with the target host architecture **s390x**.

3. Download the OpenShift Container Platform client (**oc**) and make it available for use by entering the following commands:

```
$ curl -k https://mirror.openshift.com/pub/openshift-  
v4/${ARCH}/clients/ocp/${OCP_VERSION}/openshift-client-linux.tar.gz -o oc.tar.gz
```

```
$ tar xzf oc.tar.gz
```

```
$ chmod +x oc
```

4. Download the OpenShift Container Platform installer and make it available for use by entering the following commands:

```
$ curl -k https://mirror.openshift.com/pub/openshift-  
v4/${ARCH}/clients/ocp/${OCP_VERSION}/openshift-install-linux.tar.gz -o openshift-install-  
linux.tar.gz
```

```
$ tar zxvf openshift-install-linux.tar.gz
```

```
$ chmod +x openshift-install
```

5. Prepare the **install-config.yaml** file:

```
apiVersion: v1
baseDomain: <domain> 1
compute:
- name: worker
  replicas: 0 2
controlPlane:
  name: master
  replicas: 1 3
metadata:
  name: <name> 4
networking: 5
  clusterNetwork:
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  machineNetwork:
  - cidr: 10.0.0.0/16 6
  networkType: OVNKubernetes
  serviceNetwork:
  - 172.30.0.0/16
platform:
  none: {}
bootstrapInPlace:
  installationDisk: /dev/disk/by-id/<disk_id> 7
pullSecret: '<pull_secret>' 8
sshKey: |
  <ssh_key> 9
```

- 1 Add the cluster domain name.
- 2 Set the **compute** replicas to **0**. This makes the control plane node schedulable.
- 3 Set the **controlPlane** replicas to **1**. In conjunction with the previous **compute** setting, this setting ensures the cluster runs on a single node.
- 4 Set the **metadata** name to the cluster name.
- 5 Set the **networking** details. OVN-Kubernetes is the only allowed network plugin type for single-node clusters.
- 6 Set the **cidr** value to match the subnet of the single-node OpenShift cluster.
- 7 Set the path to the installation disk drive, for example, **/dev/disk/by-id/wwn-0x64cd98f04fde100024684cf3034da5c2**.
- 8 Copy the [pull secret from Red Hat OpenShift Cluster Manager](#) and add the contents to this configuration setting.
- 9

Add the public SSH key from the administration host so that you can log in to the cluster after installation.


6. Generate OpenShift Container Platform assets by running the following commands:

```
$ mkdir ocp
```

```
$ cp install-config.yaml ocp
```

```
$ ./openshift-install --dir=ocp create single-node-ignition-config
```

7. Obtain the RHEL **kernel**, **initramfs**, and **rootfs** artifacts from the [Product Downloads](#) page on the Red Hat Customer Portal or from the [RHCOS image mirror](#) page.



IMPORTANT

The RHCOS images might not change with every release of OpenShift Container Platform. You must download images with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Only use the appropriate **kernel**, **initramfs**, and **rootfs** artifacts described in the following procedure.

The file names contain the OpenShift Container Platform version number. They resemble the following examples:

kernel


```
rhcos-<version>-live-kernel-<architecture>
```

initramfs

```
rhcos-<version>-live-initramfs.<architecture>.img
```

rootfs

```
rhcos-<version>-live-rootfs.<architecture>.img
```



NOTE

The **rootfs** image is the same for FCP and DASD.

8. Move the following artifacts and files to an HTTP or HTTPS server:
 - Downloaded RHEL live **kernel**, **initramfs**, and **rootfs** artifacts
 - Ignition files
9. Create parameter files for a particular virtual machine:

Example parameter file

```
rd.neednet=1 \
console=ttysclp0 \
coreos.live.rootfs_url=<rhcos_liveos>:8080/rootfs.img \ 1
ignition.firstboot ignition.platform.id=metal \
```

```

ignition.config.url=<rhcos_ign>:8080/ignition/bootstrap-in-place-for-live-iso.ign \ 2
ip=encbdd0:dhcp::02:00:00:02:34:02 3
rd.znet=qeth,0.0.bdd0,0.0.bdd1,0.0.bdd2,layer2=1 \
rd.dasd=0.0.4411 \ 4
rd.zfcp=0.0.8001,0x50050763040051e3,0x4000406300000000 \ 5
zfcp.allow_lun_scan=0 \
rd.luks.options=discard

```

- 1 For the **coreos.live.rootfs_url=** artifact, specify the matching **rootfs** artifact for the **kernel** and **initramfs** you are booting. Only HTTP and HTTPS protocols are supported.
- 2 For the **ignition.config.url=** parameter, specify the Ignition file for the machine role. Only HTTP and HTTPS protocols are supported.
- 3 For the **ip=** parameter, assign the IP address automatically using DHCP or manually as described in "Installing a cluster with z/VM on IBM Z® and IBM® LinuxONE".
- 4 For installations on DASD-type disks, use **rd.dasd=** to specify the DASD where RHCOS is to be installed. Omit this entry for FCP-type disks.
- 5 For installations on FCP-type disks, use **rd.zfcp=<adapter>,<wwpn>,<lun>** to specify the FCP disk where RHCOS is to be installed. Omit this entry for DASD-type disks.

Leave all other parameters unchanged.

10. Transfer the following artifacts, files, and images to z/VM. For example by using FTP:

- **kernel** and **initramfs** artifacts
 - Parameter files
 - RHCOS images
- For details about how to transfer the files with FTP and boot from the virtual reader, see [Installing under Z/VM](#).

11. Punch the files to the virtual reader of the z/VM guest virtual machine that is to become your bootstrap node.
12. Log in to CMS on the bootstrap machine.
13. IPL the bootstrap machine from the reader by running the following command:

```
$ cp ipl c
```

14. After the first reboot of the virtual machine, run the following commands directly after one another:
 - a. To boot a DASD device after first reboot, run the following commands:

```
$ cp i <devno> clear loadparm prompt
```

where:

<devno>

Specifies the device number of the boot device as seen by the guest.

```
$ cp vi vmsg 0 <kernel_parameters>
```

where:

<kernel_parameters>

Specifies a set of kernel parameters to be stored as system control program data (SCPDATA). When booting Linux, these kernel parameters are concatenated to the end of the existing kernel parameters that are used by your boot configuration. The combined parameter string must not exceed 896 characters.

- b. To boot an FCP device after first reboot, run the following commands:

```
$ cp set loaddev portname <wwpn> lun <lun>
```

where:

<wwpn>

Specifies the target port and **<lun>** the logical unit in hexadecimal format.

```
$ cp set loaddev bootprog <n>
```

where:

<n>

Specifies the kernel to be booted.

```
$ cp set loaddev scpdata {APPEND|NEW} '<kernel_parameters>'
```

where:

<kernel_parameters>

Specifies a set of kernel parameters to be stored as system control program data (SCPDATA). When booting Linux, these kernel parameters are concatenated to the end of the existing kernel parameters that are used by your boot configuration. The combined parameter string must not exceed 896 characters.

<APPEND|NEW>

Optional: Specify **APPEND** to append kernel parameters to existing SCPDATA. This is the default. Specify **NEW** to replace existing SCPDATA.

Example

```
$ cp set loaddev scpdata  
'rd.zfcp=0.0.8001,0x500507630a0350a4,0x4000409D00000000  
ip=enbddd0:dhcp::02:00:00:02:34:02 rd.neednet=1'
```

To start the IPL and boot process, run the following command:

```
$ cp i <devno>
```

where:

<devno>

Specifies the device number of the boot device as seen by the guest.

14.2.7.2. Installing single-node OpenShift with RHEL KVM on IBM Z and IBM LinuxONE

Prerequisites

- You have installed **podman**.

Procedure

1. Set the OpenShift Container Platform version by running the following command:

```
$ OCP_VERSION=<ocp_version> 1
```

- 1 Replace **<ocp_version>** with the current version. For example, **latest-4.16**.

2. Set the host architecture by running the following command:

```
$ ARCH=<architecture> 1
```

- 1 Replace **<architecture>** with the target host architecture **s390x**.

3. Download the OpenShift Container Platform client (**oc**) and make it available for use by entering the following commands:

```
$ curl -k https://mirror.openshift.com/pub/openshift-  
v4/${ARCH}/clients/ocp/${OCP_VERSION}/openshift-client-linux.tar.gz -o oc.tar.gz
```

```
$ tar zxf oc.tar.gz
```

```
$ chmod +x oc
```

4. Download the OpenShift Container Platform installer and make it available for use by entering the following commands:

```
$ curl -k https://mirror.openshift.com/pub/openshift-  
v4/${ARCH}/clients/ocp/${OCP_VERSION}/openshift-install-linux.tar.gz -o openshift-install-  
linux.tar.gz
```

```
$ tar zxvf openshift-install-linux.tar.gz
```

```
$ chmod +x openshift-install
```

5. Prepare the **install-config.yaml** file:

```
apiVersion: v1  
baseDomain: <domain> 1  
compute:  
- name: worker
```

```

  replicas: 0 2
controlPlane:
  name: master
  replicas: 1 3
metadata:
  name: <name> 4
networking: 5
  clusterNetwork:
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  machineNetwork:
  - cidr: 10.0.0.0/16 6
  networkType: OVNKubernetes
  serviceNetwork:
  - 172.30.0.0/16
platform:
  none: {}
bootstrapInPlace:
  installationDisk: /dev/disk/by-id/<disk_id> 7
pullSecret: '<pull_secret>' 8
sshKey: |
  <ssh_key> 9

```

- 1** Add the cluster domain name.
- 2** Set the **compute** replicas to **0**. This makes the control plane node schedulable.
- 3** Set the **controlPlane** replicas to **1**. In conjunction with the previous **compute** setting, this setting ensures the cluster runs on a single node.
- 4** Set the **metadata** name to the cluster name.
- 5** Set the **networking** details. OVN-Kubernetes is the only allowed network plugin type for single-node clusters.
- 6** Set the **cidr** value to match the subnet of the single-node OpenShift cluster.
- 7** Set the path to the installation disk drive, for example, **/dev/disk/by-id/wwn-0x64cd98f04fde100024684cf3034da5c2**.
- 8** Copy the [pull secret from Red Hat OpenShift Cluster Manager](#) and add the contents to this configuration setting.
- 9** Add the public SSH key from the administration host so that you can log in to the cluster after installation.


6. Generate OpenShift Container Platform assets by running the following commands:

```
$ mkdir ocp
```

```
$ cp install-config.yaml ocp
```

```
$ ./openshift-install --dir=ocp create single-node-ignition-config
```


- Obtain the RHEL **kernel**, **initramfs**, and **rootfs** artifacts from the [Product Downloads](#) page on the Red Hat Customer Portal or from the [RHCOS image mirror](#) page.



IMPORTANT

The RHCOS images might not change with every release of OpenShift Container Platform. You must download images with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Only use the appropriate **kernel**, **initramfs**, and **rootfs** artifacts described in the following procedure.

The file names contain the OpenShift Container Platform version number. They resemble the following examples:

kernel

rhcos-<version>-live-kernel-<architecture>

initramfs

rhcos-<version>-live-initramfs.<architecture>.img

rootfs

rhcos-<version>-live-rootfs.<architecture>.img

- Before you launch **virt-install**, move the following files and artifacts to an HTTP or HTTPS server:
 - Downloaded RHEL live **kernel**, **initramfs**, and **rootfs** artifacts
 - Ignition files
- Create the KVM guest nodes by using the following components:
 - RHEL **kernel** and **initramfs** artifacts
 - Ignition files
 - The new disk image
 - Adjusted parm line arguments

```
$ virt-install \
  --name <vm_name> \
  --autostart \
  --memory=<memory_mb> \
  --cpu host \
  --vcpus <vcpus> \
  --location <media_location>,kernel=<rhcos_kernel>,initrd=<rhcos_initrd> \ 1
  --disk size=100 \
  --network network=<virt_network_parm> \
  --graphics none \
  --noautoconsole \
  --extra-args "ip=<ip>::<gateway>:<mask>:<hostname>::none" \
  --extra-args "nameserver=<name_server>" \
  --extra-args "ip=dhcp rd.neednet=1 ignition.platform.id=metal ignition.firstboot" \
  --extra-args "coreos.live.rootfs_url=<rhcos_liveos>" \ 2
  --extra-args "ignition.config.url=<rhcos_ign>" \ 3
```

```
--extra-args "random.trust_cpu=on rd.luks.options=discard" \
--extra-args "console=ttysclp0" \
--wait
```

- 1 For the **--location** parameter, specify the location of the kernel/initrd on the HTTP or HTTPS server.
- 2 For the **coreos.live.rootfs_url=** artifact, specify the matching **rootfs** artifact for the **kernel** and **initramfs** you are booting. Only HTTP and HTTPS protocols are supported.
- 3 For the **ignition.config.url=** parameter, specify the Ignition file for the machine role. Only HTTP and HTTPS protocols are supported.

14.2.7.3. Installing single-node OpenShift in an LPAR on IBM Z and IBM LinuxONE

Prerequisites

- If you are deploying a single-node cluster there are zero compute nodes, the Ingress Controller pods run on the control plane nodes. In single-node cluster deployments, you must configure your application ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes. See the *Load balancing requirements for user-provisioned infrastructure* section for more information.

Procedure

1. Set the OpenShift Container Platform version by running the following command:

```
$ OCP_VERSION=<ocp_version> 1
```

- 1 Replace **<ocp_version>** with the current version. For example, **latest-4.16**.

2. Set the host architecture by running the following command:

```
$ ARCH=<architecture> 1
```

- 1 Replace **<architecture>** with the target host architecture **s390x**.

3. Download the OpenShift Container Platform client (**oc**) and make it available for use by entering the following commands:

```
$ curl -k https://mirror.openshift.com/pub/openshift-
v4/${ARCH}/clients/ocp/${OCP_VERSION}/openshift-client-linux.tar.gz -o oc.tar.gz
```

```
$ tar zxvf oc.tar.gz
```

```
$ chmod +x oc
```

4. Download the OpenShift Container Platform installer and make it available for use by entering the following commands:

```
$ curl -k https://mirror.openshift.com/pub/openshift-
v4/${ARCH}/clients/ocp/${OCP_VERSION}/openshift-install-linux.tar.gz -o openshift-install-
linux.tar.gz
```

```
$ tar zxvf openshift-install-linux.tar.gz
```

```
$ chmod +x openshift-install
```

5. Prepare the **install-config.yaml** file:

```
apiVersion: v1
baseDomain: <domain> 1
compute:
- name: worker
  replicas: 0 2
controlPlane:
  name: master
  replicas: 1 3
metadata:
  name: <name> 4
networking: 5
  clusterNetwork:
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  machineNetwork:
  - cidr: 10.0.0.0/16 6
  networkType: OVNKubernetes
  serviceNetwork:
  - 172.30.0.0/16
platform:
  none: {}
pullSecret: '<pull_secret>' 7
sshKey: |
  <ssh_key> 8
```

- 1 Add the cluster domain name.
- 2 Set the **compute** replicas to **0**. This makes the control plane node schedulable.
- 3 Set the **controlPlane** replicas to **1**. In conjunction with the previous **compute** setting, this setting ensures the cluster runs on a single node.
- 4 Set the **metadata** name to the cluster name.
- 5 Set the **networking** details. OVN-Kubernetes is the only allowed network plugin type for single-node clusters.
- 6 Set the **cidr** value to match the subnet of the single-node OpenShift cluster.
- 7 Copy the [pull secret from Red Hat OpenShift Cluster Manager](#) and add the contents to this configuration setting.
- 8 Add the public SSH key from the administration host so that you can log in to the cluster after installation.

6. Generate OpenShift Container Platform assets by running the following commands:

```
$ mkdir ocp
```

```
$ cp install-config.yaml ocp
```

7. Change to the directory that contains the OpenShift Container Platform installation program and generate the Kubernetes manifests for the cluster:

```
$ ./openshift-install create manifests --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the installation directory that contains the **install-config.yaml** file you created.

8. Check that the **mastersSchedulable** parameter in the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** Kubernetes manifest file is set to **true**.

- a. Open the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** file.
- b. Locate the **mastersSchedulable** parameter and ensure that it is set to **true** as shown in the following **spec** stanza:

```
spec:
  mastersSchedulable: true
status: {}
```

- c. Save and exit the file.

9. Create the Ignition configuration files by running the following command from the directory that contains the installation program:

```
$ ./openshift-install create ignition-configs --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the same installation directory.

10. Obtain the RHEL **kernel**, **initramfs**, and **rootfs** artifacts from the [Product Downloads](#) page on the Red Hat Customer Portal or from the [RHCOS image mirror](#) page.



IMPORTANT

The RHCOS images might not change with every release of OpenShift Container Platform. You must download images with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Only use the appropriate **kernel**, **initramfs**, and **rootfs** artifacts described in the following procedure.

The file names contain the OpenShift Container Platform version number. They resemble the following examples:

kernel


rhcos-<version>-live-kernel-<architecture>

initramfs

rhcos-<version>-live-initramfs.<architecture>.img

rootfs

rhcos-<version>-live-rootfs.<architecture>.img



NOTE

The **rootfs** image is the same for FCP and DASD.

11. Move the following artifacts and files to an HTTP or HTTPS server:
 - Downloaded RHEL live **kernel**, **initramfs**, and **rootfs** artifacts
 - Ignition files
12. Create a parameter file for the bootstrap in an LPAR:

Example parameter file for the bootstrap machine

```

cio_ignore=all,!condev rd.neednet=1 \
console=ttySCLP0 \
coreos.inst.install_dev=/dev/<block_device> \ 1
coreos.inst.ignition_url=http://<http_server>/bootstrap.ign \ 2
coreos.live.rootfs_url=http://<http_server>/rhcos-<version>-live-rootfs.<architecture>.img \ 3
ip=<ip>::<gateway>:<netmask>:<hostname>::none nameserver=<dns> \ 4
rd.znet=qeth,0.0.1140,0.0.1141,0.0.1142,layer2=1,portno=0 \
rd.dasd=0.0.4411 \ 5
rd.zfcp=0.0.8001,0x50050763040051e3,0x4000406300000000 \ 6
zfcp.allow_lun_scan=0

```

- 1 Specify the block device on the system to install to. For installations on DASD-type disk use **dasda**, for installations on FCP-type disks use **sda**.
- 2 Specify the location of the **bootstrap.ign** config file. Only HTTP and HTTPS protocols are supported.
- 3 For the **coreos.live.rootfs_url=** artifact, specify the matching **rootfs** artifact for the **kernel** and **initramfs** you are booting. Only HTTP and HTTPS protocols are supported.
- 4 For the **ip=** parameter, assign the IP address manually as described in "Installing a cluster in an LPAR on IBM Z® and IBM® LinuxONE".
- 5 For installations on DASD-type disks, use **rd.dasd=** to specify the DASD where RHCOS is to be installed. Omit this entry for FCP-type disks.
- 6 For installations on FCP-type disks, use **rd.zfcp=<adapter>,<wwpn>,<lun>** to specify the FCP disk where RHCOS is to be installed. Omit this entry for DASD-type disks.

You can adjust further parameters if required.

13. Create a parameter file for the control plane in an LPAR:

Example parameter file for the control plane machine

```

cio_ignore=all,!condev rd.neednet=1 \
console=ttySCLP0 \
coreos.inst.install_dev=/dev/<block_device> \
coreos.inst.ignition_url=http://<http_server>/master.ign 1 \
coreos.live.rootfs_url=http://<http_server>/rhcos-<version>-live-rootfs.<architecture>.img \
ip=<ip>:<gateway>:<netmask>:<hostname>::none nameserver=<dns> \
rd.znet=qeth,0.0.1140,0.0.1141,0.0.1142,layer2=1,portno=0 \
rd.dasd=0.0.4411 \
rd.zfcp=0.0.8001,0x50050763040051e3,0x4000406300000000 \
zfcp.allow_lun_scan=0

```

- 1** Specify the location of the **master.ign** config file. Only HTTP and HTTPS protocols are supported.

14. Transfer the following artifacts, files, and images to the LPAR. For example by using FTP:

- **kernel** and **initramfs** artifacts
- Parameter files
- RHCOS images

For details about how to transfer the files with FTP and boot, see [Installing in an LPAR](#).

15. Boot the bootstrap machine.

16. Boot the control plane machine.

14.2.8. Installing single-node OpenShift with IBM Power

Installing a single-node cluster on IBM Power® requires user-provisioned installation using the "Installing a cluster with IBM Power®" procedure.



NOTE

Installing a single-node cluster on IBM Power® simplifies installation for development and test environments and requires less resource requirements at entry level.

Hardware requirements

- The equivalent of two Integrated Facilities for Linux (IFL), which are SMT2 enabled, for each cluster.
- At least one network connection to connect to the **LoadBalancer** service and to serve data for traffic outside of the cluster.



NOTE

You can use dedicated or shared IFLs to assign sufficient compute resources. Resource sharing is one of the key strengths of IBM Power®. However, you must adjust capacity correctly on each hypervisor layer and ensure sufficient resources for every OpenShift Container Platform cluster.

Additional resources

- [Installing a cluster on IBM Power®](#)

14.2.8.1. Setting up basion for single-node OpenShift with IBM Power

Prior to installing single-node OpenShift on IBM Power®, you must set up bastion. Setting up a bastion server for single-node OpenShift on IBM Power® requires the configuration of the following services:

- PXE is used for the single-node OpenShift cluster installation. PXE requires the following services to be configured and run:
 - DNS to define api, api-int, and *.apps
 - DHCP service to enable PXE and assign an IP address to single-node OpenShift node
 - HTTP to provide ignition and RHCOS rootfs image
 - TFTP to enable PXE
- You must install **dnsmasq** to support DNS, DHCP and PXE, httpd for HTTP.

Use the following procedure to configure a bastion server that meets these requirements.

Procedure

1. Use the following command to install **grub2**, which is required to enable PXE for PowerVM:

```
grub2-mknetdir --net-directory=/var/lib/tftpboot
```

Example `/var/lib/tftpboot/boot/grub2/grub.cfg` file

```
default=0
fallback=1
timeout=1
if [ ${net_default_mac} == fa:b0:45:27:43:20 ]; then
menuentry "CoreOS (BIOS)" {
    echo "Loading kernel"
    linux "/rhcos/kernel" ip=dhcp rd.neednet=1 ignition.platform.id=metal ignition.firstboot
coreos.live.rootfs_url=http://192.168.10.5:8000/install/rootfs.img
    ignition.config.url=http://192.168.10.5:8000/ignition/sno.ign
    echo "Loading initrd"
    initrd "/rhcos/initramfs.img"
}
fi
```

2. Use the following commands to download RHCOS image files from the mirror repo for PXE.
 - a. Enter the following command to assign the **RHCOS_URL** variable the follow 4.12 URL:

```
$ export RHCOS_URL=https://mirror.openshift.com/pub/openshift-
v4/ppc64le/dependencies/rhcos/4.12/latest/
```

- b. Enter the following command to navigate to the `/var/lib/tftpboot/rhcos` directory:

```
$ cd /var/lib/tftpboot/rhcos
```

- c. Enter the following command to download the specified RHCOS kernel file from the URL stored in the **RHCOS_URL** variable:

```
$ wget ${RHCOS_URL}/rhcos-live-kernel-ppc64le -o kernel
```

- d. Enter the following command to download the RHCOS **initramfs** file from the URL stored in the **RHCOS_URL** variable:

```
$ wget ${RHCOS_URL}/rhcos-live-initramfs.ppc64le.img -o initramfs.img
```

- e. Enter the following command to navigate to the **/var//var/www/html/install/** directory:

```
$ cd /var//var/www/html/install/
```

- f. Enter the following command to download, and save, the RHCOS **root filesystem** image file from the URL stored in the **RHCOS_URL** variable:

```
$ wget ${RHCOS_URL}/rhcos-live-rootfs.ppc64le.img -o rootfs.img
```

3. To create the ignition file for a single-node OpenShift cluster, you must create the **install-config.yaml** file.

- a. Enter the following command to create the work directory that holds the file:

```
$ mkdir -p ~/sno-work
```

- b. Enter the following command to navigate to the **~/sno-work** directory:

```
$ cd ~/sno-work
```

- c. Use the following sample file can to create the required **install-config.yaml** in the **~/sno-work** directory:

```
apiVersion: v1
baseDomain: <domain> 1
compute:
- name: worker
  replicas: 0 2
controlPlane:
  name: master
  replicas: 1 3
metadata:
  name: <name> 4
networking: 5
  clusterNetwork:
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  machineNetwork:
  - cidr: 10.0.0.0/16 6
  networkType: OVNKubernetes
```



```

serviceNetwork:
  - 172.30.0.0/16
platform:
  none: {}
bootstrapInPlace:
  installationDisk: /dev/disk/by-id/<disk_id> 7
pullSecret: '<pull_secret>' 8
sshKey: |
  <ssh_key> 9

```

- 1 Add the cluster domain name.
 - 2 Set the **compute** replicas to **0**. This makes the control plane node schedulable.
 - 3 Set the **controlPlane** replicas to **1**. In conjunction with the previous **compute** setting, this setting ensures that the cluster runs on a single node.
 - 4 Set the **metadata** name to the cluster name.
 - 5 Set the **networking** details. OVN-Kubernetes is the only allowed network plugin type for single-node clusters.
 - 6 Set the **cidr** value to match the subnet of the single-node OpenShift cluster.
 - 7 Set the path to the installation disk drive, for example, **/dev/disk/by-id/wwn-0x64cd98f04fde100024684cf3034da5c2**.
 - 8 Copy the [pull secret from Red Hat OpenShift Cluster Manager](#) and add the contents to this configuration setting.
 - 9 Add the public SSH key from the administration host so that you can log in to the cluster after installation.
4. Download the **openshift-install** image to create the ignition file and copy it to the **http** directory.
- a. Enter the following command to download the **openshift-install-linux-4.12.0** .tar file:


```
$ wget https://mirror.openshift.com/pub/openshift-
v4/ppc64le/clients/ocp/4.12.0/openshift-install-linux-4.12.0.tar.gz
```
 - b. Enter the following command to unpack the **openshift-install-linux-4.12.0.tar.gz** archive:


```
$ tar xzvf openshift-install-linux-4.12.0.tar.gz
```
 - c. Enter the following command to


```
$ ./openshift-install --dir=~/.sno-work create create single-node-ignition-config
```
 - d. Enter the following command to create the ignition file:


```
$ cp ~/.sno-work/single-node-ignition-config.ign /var/www/html/ignition/sno.ign
```
 - e. Enter the following command to restore SELinux file for the **/var/www/html** directory:

```
$ restorecon -vR /var/www/html || true
```

Bastion now has all the required files and is properly configured in order to install single-node OpenShift.

14.2.8.2. Installing single-node OpenShift with IBM Power

Prerequisites

- You have set up bastion.

Procedure

There are two steps for the single-node OpenShift cluster installation. First the single-node OpenShift logical partition (LPAR) needs to boot up with PXE, then you need to monitor the installation progress.

1. Use the following command to boot powerVM with netboot:

```
$ lpar_netboot -i -D -f -t ent -m <sno_mac> -s auto -d auto -S <server_ip> -C <sno_ip> -G <gateway> <lpar_name> default_profile <cec_name>
```

where:

sno_mac

Specifies the MAC address of the single-node OpenShift cluster.

sno_ip

Specifies the IP address of the single-node OpenShift cluster.

server_ip

Specifies the IP address of bastion (PXE server).

gateway

Specifies the Network's gateway IP.

lpar_name

Specifies the single-node OpenShift lpar name in HMC.

cec_name

Specifies the System name where the sno_lpar resides

2. After the single-node OpenShift LPAR boots up with PXE, use the **openshift-install** command to monitor the progress of installation:

- a. Run the following command after the bootstrap is complete:

```
./openshift-install wait-for bootstrap-complete
```

- b. Run the following command after it returns successfully:

```
./openshift-install wait-for install-complete
```

CHAPTER 15. INSTALLING ON BARE METAL

15.1. PREPARING FOR BARE METAL CLUSTER INSTALLATION

15.1.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You have read the documentation on [selecting a cluster installation method and preparing it for users](#).

15.1.2. Planning a bare metal cluster for OpenShift Virtualization

If you will use OpenShift Virtualization, it is important to be aware of several requirements before you install your bare metal cluster.

- If you want to use live migration features, you must have multiple worker nodes *at the time of cluster installation*. This is because live migration requires the cluster-level high availability (HA) flag to be set to true. The HA flag is set when a cluster is installed and cannot be changed afterwards. If there are fewer than two worker nodes defined when you install your cluster, the HA flag is set to false for the life of the cluster.



NOTE

You can install OpenShift Virtualization on a single-node cluster, but single-node OpenShift does not support high availability.

- Live migration requires shared storage. Storage for OpenShift Virtualization must support and use the ReadWriteMany (RWX) access mode.
- If you plan to use Single Root I/O Virtualization (SR-IOV), ensure that your network interface controllers (NICs) are supported by OpenShift Container Platform.

Additional resources

- [Getting started with OpenShift Virtualization](#)
- [Preparing your cluster for OpenShift Virtualization](#)
- [About Single Root I/O Virtualization \(SR-IOV\) hardware networks](#)
- [Connecting a virtual machine to an SR-IOV network](#)

15.1.3. NIC partitioning for SR-IOV devices (Technology Preview)

OpenShift Container Platform can be deployed on a server with a dual port network interface card (NIC). You can partition a single, high-speed dual port NIC into multiple virtual functions (VFs) and enable SR-IOV.

This feature supports the use of bonds for high availability with the Link Aggregation Control Protocol (LACP).

**NOTE**

Only one LACP can be declared by physical NIC.

An OpenShift Container Platform cluster can be deployed on a bond interface with 2 VFs on 2 physical functions (PFs) using the following methods:


- Agent-based installer

**NOTE**

The minimum required version of **nmstate** is:

- **1.4.2-4** for RHEL 8 versions
- **2.2.7** for RHEL 9 versions

- Installer-provisioned infrastructure installation
- User-provisioned infrastructure installation

**IMPORTANT**

Support for Day 1 operations associated with enabling NIC partitioning for SR-IOV devices is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

Additional resources

- [Example: Bonds and SR-IOV dual-nic node network configuration](#)
- [Optional: Configuring host network interfaces for dual port NIC](#)
- [Bonding multiple SR-IOV network interfaces to a dual port NIC interface](#)

15.1.4. Choosing a method to install OpenShift Container Platform on bare metal

The OpenShift Container Platform installation program offers four methods for deploying a cluster:

- **Interactive:** You can deploy a cluster with the web-based [Assisted Installer](#). This is the recommended approach for clusters with networks connected to the internet. The Assisted Installer is the easiest way to install OpenShift Container Platform, it provides smart defaults, and it performs pre-flight validations before installing the cluster. It also provides a RESTful API for automation and advanced configuration scenarios.
- **Local Agent-based:** You can deploy a cluster locally with the [agent-based installer](#) for air-gapped or restricted networks. It provides many of the benefits of the Assisted Installer, but you must download and configure the [agent-based installer](#) first. Configuration is done with a commandline interface. This approach is ideal for air-gapped or restricted networks.

- **Automated:** You can [deploy a cluster on installer-provisioned infrastructure](#) and the cluster it maintains. The installer uses each cluster host's baseboard management controller (BMC) for provisioning. You can deploy clusters with both connected or air-gapped or restricted networks.
- **Full control:** You can deploy a cluster on [infrastructure that you prepare and maintain](#), which provides maximum customizability. You can deploy clusters with both connected or air-gapped or restricted networks.

The clusters have the following characteristics:

- Highly available infrastructure with no single points of failure is available by default.
- Administrators maintain control over what updates are applied and when.

See [Installation process](#) for more information about installer-provisioned and user-provisioned installation processes.

15.1.4.1. Installing a cluster on installer-provisioned infrastructure

You can install a cluster on bare metal infrastructure that is provisioned by the OpenShift Container Platform installation program, by using the following method:

- [Installing an installer-provisioned cluster on bare metal](#) You can install OpenShift Container Platform on bare metal by using installer provisioning.


15.1.4.2. Installing a cluster on user-provisioned infrastructure

You can install a cluster on bare metal infrastructure that you provision, by using one of the following methods:

- [Installing a user-provisioned cluster on bare metal](#) You can install OpenShift Container Platform on bare metal infrastructure that you provision. For a cluster that contains user-provisioned infrastructure, you must deploy all of the required machines.
- [Installing a user-provisioned bare metal cluster with network customizations](#) You can install a bare metal cluster on user-provisioned infrastructure with network-customizations. By customizing your network configuration, your cluster can coexist with existing IP address allocations in your environment and integrate with existing MTU and VXLAN configurations. Most of the network customizations must be applied at the installation stage.
- [Installing a user-provisioned bare metal cluster on a restricted network](#) You can install a user-provisioned bare metal cluster on a restricted or disconnected network by using a mirror registry. You can also use this installation method to ensure that your clusters only use container images that satisfy your organizational controls on external content.

15.2. INSTALLING A USER-PROVISIONED CLUSTER ON BARE METAL

In OpenShift Container Platform 4.16, you can install a cluster on bare metal infrastructure that you provision.



IMPORTANT

While you might be able to follow this procedure to deploy a cluster on virtualized or cloud environments, you must be aware of additional considerations for non-bare metal platforms. Review the information in the [guidelines for deploying OpenShift Container Platform on non-tested platforms](#) before you attempt to install an OpenShift Container Platform cluster in such an environment.

15.2.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.



NOTE

Be sure to also review this site list if you are configuring a proxy.

15.2.2. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.



IMPORTANT

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

Additional resources

- See [Installing a user-provisioned bare metal cluster on a restricted network](#) for more information about performing a restricted network installation on bare metal infrastructure that you provision.

15.2.3. Requirements for a cluster with user-provisioned infrastructure

For a cluster that contains user-provisioned infrastructure, you must deploy all of the required machines.

This section describes the requirements for deploying OpenShift Container Platform on user-provisioned infrastructure.

15.2.3.1. Required machines for cluster installation

The smallest OpenShift Container Platform clusters require the following hosts:

Table 15.1. Minimum required hosts

Hosts	Description
One temporary bootstrap machine	The cluster requires the bootstrap machine to deploy the OpenShift Container Platform cluster on the three control plane machines. You can remove the bootstrap machine after you install the cluster.
Three control plane machines	The control plane machines run the Kubernetes and OpenShift Container Platform services that form the control plane.
At least two compute machines, which are also known as worker machines.	The workloads requested by OpenShift Container Platform users run on the compute machines.



NOTE

As an exception, you can run zero compute machines in a bare metal cluster that consists of three control plane machines only. This provides smaller, more resource efficient clusters for cluster administrators and developers to use for testing, development, and production. Running one compute machine is not supported.



IMPORTANT

To maintain high availability of your cluster, use separate physical hosts for these cluster machines.

The bootstrap and control plane machines must use Red Hat Enterprise Linux CoreOS (RHCOS) as the operating system. However, the compute machines can choose between Red Hat Enterprise Linux CoreOS (RHCOS), Red Hat Enterprise Linux (RHEL) 8.6 and later.

Note that RHCOS is based on Red Hat Enterprise Linux (RHEL) 9.2 and inherits all of its hardware certifications and requirements. See [Red Hat Enterprise Linux technology capabilities and limits](#) .

15.2.3.2. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 15.2. Minimum resource requirements

Machine	Operating System	CPU [1]	RAM	Storage	Input/Output Per Second (IOPS)[2]
Bootstrap	RHCOS	4	16 GB	100 GB	300
Control plane	RHCOS	4	16 GB	100 GB	300
Compute	RHCOS, RHEL 8.6 and later [3]	2	8 GB	100 GB	300

1. One CPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio: (threads per core × cores) × sockets = CPUs.
2. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes which require a 10 ms p99 fsync duration. Note that on many cloud platforms, storage size and IOPS scale together, so you might need to over-allocate storage volume to obtain sufficient performance.
3. As with all user-provisioned installations, if you choose to use RHEL compute machines in your cluster, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Use of RHEL 7 compute machines is deprecated and has been removed in OpenShift Container Platform 4.10 and later.

NOTE

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

15.2.3.3. Certificate signing requests management

Because your cluster has limited access to automatic machine management when you use infrastructure that you provision, you must provide a mechanism for approving cluster certificate signing requests (CSRs) after installation. The **kube-controller-manager** only approves the kubelet client CSRs. The **machine-approver** cannot guarantee the validity of a serving certificate that is requested by using kubelet credentials because it cannot confirm that the correct machine issued the request. You must determine and implement a method of verifying the validity of the kubelet serving certificate requests and approving them.

Additional resources

- See [Configuring a three-node cluster](#) for details about deploying three-node clusters in bare metal environments.
- See [Approving the certificate signing requests for your machines](#) for more information about approving cluster certificate signing requests after installation.

15.2.3.4. Requirements for baremetal clusters on vSphere

Ensure you enable the **disk.EnableUUID** parameter on all virtual machines in your cluster.

Additional resources

- See [Installing RHCOS and starting the OpenShift Container Platform bootstrap process](#) for details on setting the **disk.EnableUUID** parameter's value to **TRUE** on VMware vSphere for user-provisioned infrastructure.

15.2.3.5. Networking requirements for user-provisioned infrastructure

All the Red Hat Enterprise Linux CoreOS (RHCOS) machines require networking to be configured in **initramfs** during boot to fetch their Ignition config files.

During the initial boot, the machines require an IP address configuration that is set either through a DHCP server or statically by providing the required boot options. After a network connection is established, the machines download their Ignition config files from an HTTP or HTTPS server. The Ignition config files are then used to set the exact state of each machine. The Machine Config Operator completes more changes to the machines, such as the application of new certificates or keys, after installation.

It is recommended to use a DHCP server for long-term management of the cluster machines. Ensure that the DHCP server is configured to provide persistent IP addresses, DNS server information, and hostnames to the cluster machines.



NOTE

If a DHCP service is not available for your user-provisioned infrastructure, you can instead provide the IP networking configuration and the address of the DNS server to the nodes at RHCOS install time. These can be passed as boot arguments if you are installing from an ISO image. See the *Installing RHCOS and starting the OpenShift Container Platform bootstrap process* section for more information about static IP provisioning and advanced networking options.

The Kubernetes API server must be able to resolve the node names of the cluster machines. If the API servers and worker nodes are in different zones, you can configure a default DNS search zone to allow the API server to resolve the node names. Another supported approach is to always refer to hosts by their fully-qualified domain names in both the node objects and all DNS requests.

15.2.3.5.1. Setting the cluster node hostnames through DHCP

On Red Hat Enterprise Linux CoreOS (RHCOS) machines, the hostname is set through NetworkManager. By default, the machines obtain their hostname through DHCP. If the hostname is not provided by DHCP, set statically through kernel arguments, or another method, it is obtained through a reverse DNS lookup. Reverse DNS lookup occurs after the network has been initialized on a node and can take time to resolve. Other system services can start prior to this and detect the hostname as **localhost** or similar. You can avoid this by using DHCP to provide the hostname for each cluster node.

Additionally, setting the hostnames through DHCP can bypass any manual DNS record name configuration errors in environments that have a DNS split-horizon implementation.

15.2.3.5.2. Network connectivity requirements

You must configure the network connectivity between machines to allow OpenShift Container Platform cluster components to communicate. Each machine must be able to resolve the hostnames of all other machines in the cluster.

This section provides details about the ports that are required.



IMPORTANT

In connected OpenShift Container Platform environments, all nodes are required to have internet access to pull images for platform containers and provide telemetry data to Red Hat.

Table 15.3. Ports used for all-machine to all-machine communications

Protocol	Port	Description
ICMP	N/A	Network reachability tests
TCP	1936	Metrics
	9000-9999	Host level services, including the node exporter on ports 9100-9101 and the Cluster Version Operator on port 9099 .
	10250-10259	The default ports that Kubernetes reserves
UDP	4789	VXLAN
	6081	Geneve
	9000-9999	Host level services, including the node exporter on ports 9100-9101 .
	500	IPsec IKE packets
	4500	IPsec NAT-T packets

Protocol	Port	Description
	123	Network Time Protocol (NTP) on UDP port 123 If an external NTP time server is configured, you must open UDP port 123 .
TCP/UDP	30000-32767	Kubernetes node port
ESP	N/A	IPsec Encapsulating Security Payload (ESP)

Table 15.4. Ports used for all-machine to control plane communications

Protocol	Port	Description
TCP	6443	Kubernetes API

Table 15.5. Ports used for control plane machine to control plane machine communications

Protocol	Port	Description
TCP	2379-2380	etcd server and peer ports

NTP configuration for user-provisioned infrastructure

OpenShift Container Platform clusters are configured to use a public Network Time Protocol (NTP) server by default. If you want to use a local enterprise NTP server, or if your cluster is being deployed in a disconnected network, you can configure the cluster to use a specific time server. For more information, see the documentation for *Configuring chrony time service*.

If a DHCP server provides NTP server information, the chrony time service on the Red Hat Enterprise Linux CoreOS (RHCOS) machines read the information and can sync the clock with the NTP servers.

Additional resources

- [Configuring chrony time service](#)

15.2.3.6. User-provisioned DNS requirements

In OpenShift Container Platform deployments, DNS name resolution is required for the following components:

- The Kubernetes API
- The OpenShift Container Platform application wildcard
- The bootstrap, control plane, and compute machines

Reverse DNS resolution is also required for the Kubernetes API, the bootstrap machine, the control plane machines, and the compute machines.

DNS A/AAAA or CNAME records are used for name resolution and PTR records are used for reverse name resolution. The reverse records are important because Red Hat Enterprise Linux CoreOS (RHCOS) uses the reverse records to set the hostnames for all the nodes, unless the hostnames are provided by DHCP. Additionally, the reverse records are used to generate the certificate signing requests (CSR) that OpenShift Container Platform needs to operate.




NOTE

It is recommended to use a DHCP server to provide the hostnames to each cluster node. See the *DHCP recommendations for user-provisioned infrastructure* section for more information.

The following DNS records are required for a user-provisioned OpenShift Container Platform cluster and they must be in place before installation. In each record, **<cluster_name>** is the cluster name and **<base_domain>** is the base domain that you specify in the **install-config.yaml** file. A complete DNS record takes the form: **<component>.<cluster_name>.<base_domain>..**

Table 15.6. Required DNS records

Component	Record	Description
Kubernetes API	api.<cluster_name>.<base_domain>.	A DNS A/AAAA or CNAME record, and a DNS PTR record, to identify the API load balancer. These records must be resolvable by both clients external to the cluster and from all the nodes within the cluster.
	api-int.<cluster_name>.<base_domain>.	A DNS A/AAAA or CNAME record, and a DNS PTR record, to internally identify the API load balancer. These records must be resolvable from all the nodes within the cluster.
		 <p>IMPORTANT</p> <p>The API server must be able to resolve the worker nodes by the hostnames that are recorded in Kubernetes. If the API server cannot resolve the node names, then proxied API calls can fail, and you cannot retrieve logs from pods.</p>
Routes	*.apps.<cluster_name>.<base_domain>.	<p>A wildcard DNS A/AAAA or CNAME record that refers to the application ingress load balancer. The application ingress load balancer targets the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default. These records must be resolvable by both clients external to the cluster and from all the nodes within the cluster.</p> <p>For example, console-openshift-console.apps.<cluster_name>.<base_domain> is used as a wildcard route to the OpenShift Container Platform console.</p>

Component	Record	Description
Bootstrap machine	bootstrap.<cluster_name>.<base_domain>.	A DNS A/AAAA or CNAME record, and a DNS PTR record, to identify the bootstrap machine. These records must be resolvable by the nodes within the cluster.
Control plane machines	<control_plane><n>.<cluster_name>.<base_domain>.	DNS A/AAAA or CNAME records and DNS PTR records to identify each machine for the control plane nodes. These records must be resolvable by the nodes within the cluster.
Compute machines	<compute><n>.<cluster_name>.<base_domain>.	DNS A/AAAA or CNAME records and DNS PTR records to identify each machine for the worker nodes. These records must be resolvable by the nodes within the cluster.

**NOTE**

In OpenShift Container Platform 4.4 and later, you do not need to specify etcd host and SRV records in your DNS configuration.

TIP

You can use the **dig** command to verify name and reverse name resolution. See the section on *Validating DNS resolution for user-provisioned infrastructure* for detailed validation steps.

15.2.3.6.1. Example DNS configuration for user-provisioned clusters

This section provides A and PTR record configuration samples that meet the DNS requirements for deploying OpenShift Container Platform on user-provisioned infrastructure. The samples are not meant to provide advice for choosing one DNS solution over another.

In the examples, the cluster name is **ocp4** and the base domain is **example.com**.

Example DNS A record configuration for a user-provisioned cluster

The following example is a BIND zone file that shows sample A records for name resolution in a user-provisioned cluster.

Example 15.1. Sample DNS zone database


```
$TTL 1W
@ IN SOA ns1.example.com. root (
  2019070700 ; serial
  3H ; refresh (3 hours)
  30M ; retry (30 minutes)
  2W ; expiry (2 weeks)
  1W ) ; minimum (1 week)
IN NS ns1.example.com.
IN MX 10 smtp.example.com.
```

```

;
;
;
ns1.example.com. IN A 192.168.1.5
smtp.example.com. IN A 192.168.1.5
;
helper.example.com. IN A 192.168.1.5
helper.ocp4.example.com. IN A 192.168.1.5
;
api.ocp4.example.com. IN A 192.168.1.5 1
api-int.ocp4.example.com. IN A 192.168.1.5 2
;
*.apps.ocp4.example.com. IN A 192.168.1.5 3
;
bootstrap.ocp4.example.com. IN A 192.168.1.96 4
;
control-plane0.ocp4.example.com. IN A 192.168.1.97 5
control-plane1.ocp4.example.com. IN A 192.168.1.98 6
control-plane2.ocp4.example.com. IN A 192.168.1.99 7
;
compute0.ocp4.example.com. IN A 192.168.1.11 8
compute1.ocp4.example.com. IN A 192.168.1.7 9
;
;EOF

```

- 1 Provides name resolution for the Kubernetes API. The record refers to the IP address of the API load balancer.
- 2 Provides name resolution for the Kubernetes API. The record refers to the IP address of the API load balancer and is used for internal cluster communications.
- 3 Provides name resolution for the wildcard routes. The record refers to the IP address of the application ingress load balancer. The application ingress load balancer targets the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.



NOTE

In the example, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

- 4 Provides name resolution for the bootstrap machine.
- 5 6 7 Provides name resolution for the control plane machines.
- 8 9 Provides name resolution for the compute machines.

Example DNS PTR record configuration for a user-provisioned cluster

The following example BIND zone file shows sample PTR records for reverse name resolution in a user-provisioned cluster.

-


Example 15.2. Sample DNS zone database for reverse records

```

$TTL 1W
@ IN SOA ns1.example.com. root (
  2019070700 ; serial
  3H ; refresh (3 hours)
  30M ; retry (30 minutes)
  2W ; expiry (2 weeks)
  1W ) ; minimum (1 week)
IN NS ns1.example.com.
;
5.1.168.192.in-addr.arpa. IN PTR api.ocp4.example.com. 1
5.1.168.192.in-addr.arpa. IN PTR api-int.ocp4.example.com. 2
;
96.1.168.192.in-addr.arpa. IN PTR bootstrap.ocp4.example.com. 3
;
97.1.168.192.in-addr.arpa. IN PTR control-plane0.ocp4.example.com. 4
98.1.168.192.in-addr.arpa. IN PTR control-plane1.ocp4.example.com. 5
99.1.168.192.in-addr.arpa. IN PTR control-plane2.ocp4.example.com. 6
;
11.1.168.192.in-addr.arpa. IN PTR compute0.ocp4.example.com. 7
7.1.168.192.in-addr.arpa. IN PTR compute1.ocp4.example.com. 8
;
;EOF

```

- 1 Provides reverse DNS resolution for the Kubernetes API. The PTR record refers to the record name of the API load balancer.
- 2 Provides reverse DNS resolution for the Kubernetes API. The PTR record refers to the record name of the API load balancer and is used for internal cluster communications.
- 3 Provides reverse DNS resolution for the bootstrap machine.
- 4 5 6 Provides reverse DNS resolution for the control plane machines.
- 7 8 Provides reverse DNS resolution for the compute machines.

**NOTE**

A PTR record is not required for the OpenShift Container Platform application wildcard.

Additional resources

- [Validating DNS resolution for user-provisioned infrastructure](#)

15.2.3.7. Load balancing requirements for user-provisioned infrastructure

Before you install OpenShift Container Platform, you must provision the API and application Ingress load balancing infrastructure. In production scenarios, you can deploy the API and application Ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

**NOTE**

If you want to deploy the API and application Ingress load balancers with a Red Hat Enterprise Linux (RHEL) instance, you must purchase the RHEL subscription separately.

The load balancing infrastructure must meet the following requirements:

1. **API load balancer:** Provides a common endpoint for users, both human and machine, to interact with and configure the platform. Configure the following conditions:
 - Layer 4 load balancing only. This can be referred to as Raw TCP or SSL Passthrough mode.
 - A stateless load balancing algorithm. The options vary based on the load balancer implementation.

**IMPORTANT**

Do not configure session persistence for an API load balancer. Configuring session persistence for a Kubernetes API server might cause performance issues from excess application traffic for your OpenShift Container Platform cluster and the Kubernetes API that runs inside the cluster.

Configure the following ports on both the front and back of the load balancers:

Table 15.7. API load balancer

Port	Back-end machines (pool members)	Internal	External	Description
6443	Bootstrap and control plane. You remove the bootstrap machine from the load balancer after the bootstrap machine initializes the cluster control plane. You must configure the /readyz endpoint for the API server health check probe.	X	X	Kubernetes API server
22623	Bootstrap and control plane. You remove the bootstrap machine from the load balancer after the bootstrap machine initializes the cluster control plane.	X		Machine config server

**NOTE**

The load balancer must be configured to take a maximum of 30 seconds from the time the API server turns off the **/readyz** endpoint to the removal of the API server instance from the pool. Within the time frame after **/readyz** returns an error or becomes healthy, the endpoint must have been removed or added. Probing every 5 or 10 seconds, with two successful requests to become healthy and three to become unhealthy, are well-tested values.

2. **Application Ingress load balancer.** Provides an ingress point for application traffic flowing in from outside the cluster. A working configuration for the Ingress router is required for an OpenShift Container Platform cluster.

Configure the following conditions:

- Layer 4 load balancing only. This can be referred to as Raw TCP or SSL Passthrough mode.
- A connection-based or session-based persistence is recommended, based on the options available and types of applications that will be hosted on the platform.

TIP

If the true IP address of the client can be seen by the application Ingress load balancer, enabling source IP-based session persistence can improve performance for applications that use end-to-end TLS encryption.

Configure the following ports on both the front and back of the load balancers:

Table 15.8. Application Ingress load balancer

Port	Back-end machines (pool members)	Internal	External	Description
443	The machines that run the Ingress Controller pods, compute, or worker, by default.	X	X	HTTPS traffic
80	The machines that run the Ingress Controller pods, compute, or worker, by default.	X	X	HTTP traffic



NOTE

If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application Ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes.

15.2.3.7.1. Example load balancer configuration for user-provisioned clusters

This section provides an example API and application Ingress load balancer configuration that meets the load balancing requirements for user-provisioned clusters. The sample is an `/etc/haproxy/haproxy.cfg` configuration for an HAProxy load balancer. The example is not meant to provide advice for choosing one load balancing solution over another.

In the example, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.



NOTE

If you are using HAProxy as a load balancer and SELinux is set to **enforcing**, you must ensure that the HAProxy service can bind to the configured TCP port by running `setsebool -P haproxy_connect_any=1`.

Example 15.3. Sample API and application Ingress load balancer configuration

```
global
log      127.0.0.1 local2
pidfile  /var/run/haproxy.pid
maxconn  4000
daemon

defaults
mode            http
log             global
option          dontlognull
option http-server-close
option          redispatch
retries         3
timeout http-request 10s
timeout queue   1m
timeout connect 10s
timeout client  1m
timeout server  1m
timeout http-keep-alive 10s
timeout check   10s
maxconn         3000
listen api-server-6443 1
bind *:6443
mode tcp
option httpchk GET /readyz HTTP/1.0
option log-health-checks
balance roundrobin
server bootstrap bootstrap.ocp4.example.com:6443 verify none check check-ssl inter 10s fall 2
rise 3 backup 2
server master0 master0.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
fall 2 rise 3
server master1 master1.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
fall 2 rise 3
server master2 master2.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
fall 2 rise 3
listen machine-config-server-22623 3
bind *:22623
mode tcp
server bootstrap bootstrap.ocp4.example.com:22623 check inter 1s backup 4
server master0 master0.ocp4.example.com:22623 check inter 1s
server master1 master1.ocp4.example.com:22623 check inter 1s
server master2 master2.ocp4.example.com:22623 check inter 1s
listen ingress-router-443 5
bind *:443
mode tcp
balance source
server compute0 compute0.ocp4.example.com:443 check inter 1s
server compute1 compute1.ocp4.example.com:443 check inter 1s
listen ingress-router-80 6
bind *:80
mode tcp
balance source
server compute0 compute0.ocp4.example.com:80 check inter 1s
server compute1 compute1.ocp4.example.com:80 check inter 1s
```

- 1 Port **6443** handles the Kubernetes API traffic and points to the control plane machines.
- 2 4 The bootstrap entries must be in place before the OpenShift Container Platform cluster installation and they must be removed after the bootstrap process is complete.
- 3 Port **22623** handles the machine config server traffic and points to the control plane machines.
- 5 Port **443** handles the HTTPS traffic and points to the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.
- 6 Port **80** handles the HTTP traffic and points to the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.



NOTE

If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application Ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes.

TIP

If you are using HAProxy as a load balancer, you can check that the **haproxy** process is listening on ports **6443**, **22623**, **443**, and **80** by running **netstat -nltp** on the HAProxy node.

15.2.4. Preparing the user-provisioned infrastructure

Before you install OpenShift Container Platform on user-provisioned infrastructure, you must prepare the underlying infrastructure.

This section provides details about the high-level steps required to set up your cluster infrastructure in preparation for an OpenShift Container Platform installation. This includes configuring IP networking and network connectivity for your cluster nodes, enabling the required ports through your firewall, and setting up the required DNS and load balancing infrastructure.

After preparation, your cluster infrastructure must meet the requirements outlined in the *Requirements for a cluster with user-provisioned infrastructure* section.

Prerequisites

- You have reviewed the [OpenShift Container Platform 4.x Tested Integrations](#) page.
- You have reviewed the infrastructure requirements detailed in the *Requirements for a cluster with user-provisioned infrastructure* section.

Procedure

1. If you are using DHCP to provide the IP networking configuration to your cluster nodes, configure your DHCP service.
 - a. Add persistent IP addresses for the nodes to your DHCP server configuration. In your configuration, match the MAC address of the relevant network interface to the intended IP address for each node.

- b. When you use DHCP to configure IP addressing for the cluster machines, the machines also obtain the DNS server information through DHCP. Define the persistent DNS server address that is used by the cluster nodes through your DHCP server configuration.

**NOTE**

If you are not using a DHCP service, you must provide the IP networking configuration and the address of the DNS server to the nodes at RHCOS install time. These can be passed as boot arguments if you are installing from an ISO image. See the *Installing RHCOS and starting the OpenShift Container Platform bootstrap process* section for more information about static IP provisioning and advanced networking options.

- c. Define the hostnames of your cluster nodes in your DHCP server configuration. See the *Setting the cluster node hostnames through DHCP* section for details about hostname considerations.

**NOTE**

If you are not using a DHCP service, the cluster nodes obtain their hostname through a reverse DNS lookup.

2. Ensure that your network infrastructure provides the required network connectivity between the cluster components. See the *Networking requirements for user-provisioned infrastructure* section for details about the requirements.
3. Configure your firewall to enable the ports required for the OpenShift Container Platform cluster components to communicate. See *Networking requirements for user-provisioned infrastructure* section for details about the ports that are required.

**IMPORTANT**

By default, port **1936** is accessible for an OpenShift Container Platform cluster, because each control plane node needs access to this port.

Avoid using the Ingress load balancer to expose this port, because doing so might result in the exposure of sensitive information, such as statistics and metrics, related to Ingress Controllers.

4. Setup the required DNS infrastructure for your cluster.
 - a. Configure DNS name resolution for the Kubernetes API, the application wildcard, the bootstrap machine, the control plane machines, and the compute machines.
 - b. Configure reverse DNS resolution for the Kubernetes API, the bootstrap machine, the control plane machines, and the compute machines.
See the *User-provisioned DNS requirements* section for more information about the OpenShift Container Platform DNS requirements.
5. Validate your DNS configuration.
 - a. From your installation node, run DNS lookups against the record names of the Kubernetes API, the wildcard routes, and the cluster nodes. Validate that the IP addresses in the responses correspond to the correct components.

- b. From your installation node, run reverse DNS lookups against the IP addresses of the load balancer and the cluster nodes. Validate that the record names in the responses correspond to the correct components.
See the *Validating DNS resolution for user-provisioned infrastructure* section for detailed DNS validation steps.
6. Provision the required API and application ingress load balancing infrastructure. See the *Load balancing requirements for user-provisioned infrastructure* section for more information about the requirements.

**NOTE**


Some load balancing solutions require the DNS name resolution for the cluster nodes to be in place before the load balancing is initialized.

Additional resources

- [Requirements for a cluster with user-provisioned infrastructure](#)
- [Installing RHCOS and starting the OpenShift Container Platform bootstrap process](#)
- [Setting the cluster node hostnames through DHCP](#)
- [Advanced RHCOS installation configuration](#)
- [Networking requirements for user-provisioned infrastructure](#)
- [User-provisioned DNS requirements](#)
- [Validating DNS resolution for user-provisioned infrastructure](#)
- [Load balancing requirements for user-provisioned infrastructure](#)

15.2.5. Validating DNS resolution for user-provisioned infrastructure

You can validate your DNS configuration before installing OpenShift Container Platform on user-provisioned infrastructure.

**IMPORTANT**

The validation steps detailed in this section must succeed before you install your cluster.

Prerequisites

- You have configured the required DNS records for your user-provisioned infrastructure.

Procedure

1. From your installation node, run DNS lookups against the record names of the Kubernetes API, the wildcard routes, and the cluster nodes. Validate that the IP addresses contained in the responses correspond to the correct components.
 - a. Perform a lookup against the Kubernetes API record name. Check that the result points to the IP address of the API load balancer:

```
$ dig +noall +answer @<nameserver_ip> api.<cluster_name>.<base_domain> 1
```

- 1 Replace **<nameserver_ip>** with the IP address of the nameserver, **<cluster_name>** with your cluster name, and **<base_domain>** with your base domain name.

Example output

```
api.ocp4.example.com. 604800 IN A 192.168.1.5
```

- b. Perform a lookup against the Kubernetes internal API record name. Check that the result points to the IP address of the API load balancer:

```
$ dig +noall +answer @<nameserver_ip> api-int.<cluster_name>.<base_domain>
```

Example output

```
api-int.ocp4.example.com. 604800 IN A 192.168.1.5
```

- c. Test an example ***.apps.<cluster_name>.<base_domain>** DNS wildcard lookup. All of the application wildcard lookups must resolve to the IP address of the application ingress load balancer:

```
$ dig +noall +answer @<nameserver_ip> random.apps.<cluster_name>.<base_domain>
```

Example output

```
random.apps.ocp4.example.com. 604800 IN A 192.168.1.5
```



NOTE

In the example outputs, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

You can replace **random** with another wildcard value. For example, you can query the route to the OpenShift Container Platform console:

```
$ dig +noall +answer @<nameserver_ip> console-openshift-console.apps.<cluster_name>.<base_domain>
```

Example output

```
console-openshift-console.apps.ocp4.example.com. 604800 IN A 192.168.1.5
```

- d. Run a lookup against the bootstrap DNS record name. Check that the result points to the IP address of the bootstrap node:

```
$ dig +noall +answer @<nameserver_ip> bootstrap.<cluster_name>.<base_domain>
```

Example output

```
bootstrap.ocp4.example.com. 604800 IN A 192.168.1.96
```

- e. Use this method to perform lookups against the DNS record names for the control plane and compute nodes. Check that the results correspond to the IP addresses of each node.
2. From your installation node, run reverse DNS lookups against the IP addresses of the load balancer and the cluster nodes. Validate that the record names contained in the responses correspond to the correct components.
 - a. Perform a reverse lookup against the IP address of the API load balancer. Check that the response includes the record names for the Kubernetes API and the Kubernetes internal API:

```
$ dig +noall +answer @<nameserver_ip> -x 192.168.1.5
```

Example output

```
5.1.168.192.in-addr.arpa. 604800 IN PTR api-int.ocp4.example.com. 1
5.1.168.192.in-addr.arpa. 604800 IN PTR api.ocp4.example.com. 2
```

- 1** Provides the record name for the Kubernetes internal API.
- 2** Provides the record name for the Kubernetes API.



NOTE

A PTR record is not required for the OpenShift Container Platform application wildcard. No validation step is needed for reverse DNS resolution against the IP address of the application ingress load balancer.

- b. Perform a reverse lookup against the IP address of the bootstrap node. Check that the result points to the DNS record name of the bootstrap node:

```
$ dig +noall +answer @<nameserver_ip> -x 192.168.1.96
```

Example output

```
96.1.168.192.in-addr.arpa. 604800 IN PTR bootstrap.ocp4.example.com.
```

- c. Use this method to perform reverse lookups against the IP addresses for the control plane and compute nodes. Check that the results correspond to the DNS record names of each node.

Additional resources

- [User-provisioned DNS requirements](#)
- [Load balancing requirements for user-provisioned infrastructure](#)

15.2.6. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.



NOTE

You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.



NOTE

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```


3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the **./openshift-install gather** command.



NOTE

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```



NOTE

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program. If you install a cluster on infrastructure that you provision, you must provide the key to the installation program.

Additional resources

- [Verifying node health](#)

15.2.7. Obtaining the installation program


Before you install OpenShift Container Platform, download the installation file on the host you are using for installation.

Prerequisites

- You have a computer that runs Linux or macOS, with 500 MB of local disk space.


Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select your infrastructure provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.



IMPORTANT

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.



IMPORTANT

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.


4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#) . This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

15.2.8. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.



IMPORTANT

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**. To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**. To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.

**NOTE**

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your PATH.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

15.2.9. Manually creating the installation configuration file

Installing the cluster requires that you manually create the installation configuration file.


Prerequisites

- You have an SSH public key on your local machine to provide to the installation program. The key will be used for SSH authentication onto your cluster nodes for debugging and disaster recovery.
- You have obtained the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure


1. Create an installation directory to store your required installation assets in:

```
$ mkdir <installation_directory>
```

**IMPORTANT**

You must create a directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

2. Customize the sample **install-config.yaml** file template that is provided and save it in the **<installation_directory>**.

**NOTE**

You must name this configuration file **install-config.yaml**.

3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.

**IMPORTANT**

The **install-config.yaml** file is consumed during the next step of the installation process. You must back it up now.

Additional resources

- [Installation configuration parameters for bare metal](#)

15.2.9.1. Sample install-config.yaml file for bare metal

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.

```

apiVersion: v1
baseDomain: example.com 1
compute: 2
- hyperthreading: Enabled 3
  name: worker
  replicas: 0 4
controlPlane: 5
  hyperthreading: Enabled 6
  name: master
  replicas: 3 7
metadata:
  name: test 8
networking:
  clusterNetwork:
  - cidr: 10.128.0.0/14 9
    hostPrefix: 23 10
  networkType: OVNKubernetes 11
  serviceNetwork: 12
  - 172.30.0.0/16
platform:
  none: {} 13
fips: false 14
pullSecret: '{"auths": ...}' 15
sshKey: 'ssh-ed25519 AAAA...' 16

```

- 1 The base domain of the cluster. All DNS records must be sub-domains of this base and include the cluster name.

2 5

The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.

- 3** **6** Specifies whether to enable or disable simultaneous multithreading (SMT), or hyperthreading. By default, SMT is enabled to increase the performance of the cores in your machines. You can disable it by setting the parameter value to **Disabled**. If you disable SMT, you must disable it in all cluster machines; this includes both control plane and compute machines.



NOTE

Simultaneous multithreading (SMT) is enabled by default. If SMT is not enabled in your BIOS settings, the **hyperthreading** parameter has no effect.



IMPORTANT

If you disable **hyperthreading**, whether in the BIOS or in the **install-config.yaml** file, ensure that your capacity planning accounts for the dramatically decreased machine performance.

- 4** You must set this value to **0** when you install OpenShift Container Platform on user-provisioned infrastructure. In installer-provisioned installations, the parameter controls the number of compute machines that the cluster creates and manages for you. In user-provisioned installations, you must manually deploy the compute machines before you finish installing the cluster.



NOTE

If you are installing a three-node cluster, do not deploy any compute machines when you install the Red Hat Enterprise Linux CoreOS (RHCOS) machines.

- 7** The number of control plane machines that you add to the cluster. Because the cluster uses these values as the number of etcd endpoints in the cluster, the value must match the number of control plane machines that you deploy.
- 8** The cluster name that you specified in your DNS records.
- 9** A block of IP addresses from which pod IP addresses are allocated. This block must not overlap with existing physical networks. These IP addresses are used for the pod network. If you need to access the pods from an external network, you must configure load balancers and routers to manage the traffic.



NOTE

Class E CIDR range is reserved for a future use. To use the Class E CIDR range, you must ensure your networking environment accepts the IP addresses within the Class E CIDR range.

- 10** The subnet prefix length to assign to each individual node. For example, if **hostPrefix** is set to **23**, then each node is assigned a **/23** subnet out of the given **cidr**, which allows for 510 ($2^{(32 - 23)} - 2$) pod IP addresses. If you are required to provide access to nodes from an external network, configure load balancers and routers to manage the traffic.

11

The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.

- 12 The IP address pool to use for service IP addresses. You can enter only one IP address pool. This block must not overlap with existing physical networks. If you need to access the services from an external network, configure load balancers and routers to manage the traffic.
- 13 You must set the platform to **none**. You cannot provide additional platform configuration variables for your platform.



IMPORTANT

Clusters that are installed with the platform type **none** are unable to use some features, such as managing compute machines with the Machine API. This limitation applies even if the compute machines that are attached to the cluster are installed on a platform that would normally support the feature. This parameter cannot be changed after installation.

- 14 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.



IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 15 The [pull secret from Red Hat OpenShift Cluster Manager](#) . This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.
- 16 The SSH public key for the **core** user in Red Hat Enterprise Linux CoreOS (RHCOS).



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

Additional resources

- See [Load balancing requirements for user-provisioned infrastructure](#) for more information on the API and application ingress load balancing requirements.

- See [Enabling cluster capabilities](#) for more information on enabling cluster capabilities that were disabled prior to installation.
- See [Optional cluster capabilities in OpenShift Container Platform 4.16](#) for more information about the features provided by each capability.

15.2.9.2. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.



NOTE

For bare metal installations, if you do not assign node IP addresses from the range that is specified in the **networking.machineNetwork[].cidr** field in the **install-config.yaml** file, you must include them in the **proxy.noProxy** field.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.


For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```


- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.



NOTE

The installation program does not support the proxy **readinessEndpoints** field.



NOTE

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.



NOTE

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

15.2.9.3. Configuring a three-node cluster

Optionally, you can deploy zero compute machines in a bare metal cluster that consists of three control plane machines only. This provides smaller, more resource efficient clusters for cluster administrators and developers to use for testing, development, and production.

In three-node OpenShift Container Platform environments, the three control plane machines are schedulable, which means that your application workloads are scheduled to run on them.

Prerequisites

- You have an existing **install-config.yaml** file.

Procedure

- Ensure that the number of compute replicas is set to **0** in your **install-config.yaml** file, as shown in the following **compute** stanza:

```
compute:
- name: worker
  platform: {}
  replicas: 0
```



NOTE

You must set the value of the **replicas** parameter for the compute machines to **0** when you install OpenShift Container Platform on user-provisioned infrastructure, regardless of the number of compute machines you are deploying. In installer-provisioned installations, the parameter controls the number of compute machines that the cluster creates and manages for you. This does not apply to user-provisioned installations, where the compute machines are deployed manually.

For three-node cluster installations, follow these next steps:

- If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes. See the *Load balancing requirements for user-provisioned infrastructure* section for more information.
- When you create the Kubernetes manifest files in the following procedure, ensure that the **mastersSchedulable** parameter in the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** file is set to **true**. This enables your application workloads to run on the control plane nodes.
- Do not deploy any compute nodes when you create the Red Hat Enterprise Linux CoreOS (RHCOS) machines.

15.2.10. Creating the Kubernetes manifest and Ignition config files

Because you must modify some cluster definition files and manually start the cluster machines, you must generate the Kubernetes manifest and Ignition config files that the cluster needs to configure the machines.

The installation configuration file transforms into the Kubernetes manifests. The manifests wrap into the Ignition configuration files, which are later used to configure the cluster machines.



IMPORTANT

- The Ignition config files that the OpenShift Container Platform installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

Prerequisites

- You obtained the OpenShift Container Platform installation program.
- You created the **install-config.yaml** installation configuration file.

Procedure

1. Change to the directory that contains the OpenShift Container Platform installation program and generate the Kubernetes manifests for the cluster:


```
$ ./openshift-install create manifests --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the installation directory that contains the **install-config.yaml** file you created.



WARNING

If you are installing a three-node cluster, skip the following step to allow the control plane nodes to be schedulable.



IMPORTANT

When you configure control plane nodes from the default unschedulable to schedulable, additional subscriptions are required. This is because control plane nodes then become compute nodes.


2. Check that the **mastersSchedulable** parameter in the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** Kubernetes manifest file is set to **false**. This setting prevents pods from being scheduled on the control plane machines:
 - a. Open the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** file.

- b. Locate the **mastersSchedulable** parameter and ensure that it is set to **false**.
 - c. Save and exit the file.
3. To create the Ignition configuration files, run the following command from the directory that contains the installation program:

```
$ ./openshift-install create ignition-configs --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the same installation directory.

Ignition config files are created for the bootstrap, control plane, and compute nodes in the installation directory. The **kubeadmin-password** and **kubeconfig** files are created in the **./<installation_directory>/auth** directory:



Additional resources

- See [Recovering from expired control plane certificates](#) for more information about recovering kubelet certificates.

15.2.11. Installing RHCOS and starting the OpenShift Container Platform bootstrap process

To install OpenShift Container Platform on bare metal infrastructure that you provision, you must install Red Hat Enterprise Linux CoreOS (RHCOS) on the machines. When you install RHCOS, you must provide the Ignition config file that was generated by the OpenShift Container Platform installation program for the type of machine you are installing. If you have configured suitable networking, DNS, and load balancing infrastructure, the OpenShift Container Platform bootstrap process begins automatically after the RHCOS machines have rebooted.

To install RHCOS on the machines, follow either the steps to use an ISO image or network PXE booting.



NOTE

The compute node deployment steps included in this installation document are RHCOS-specific. If you choose instead to deploy RHEL-based compute nodes, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Only RHEL 8 compute machines are supported.

You can configure RHCOS during ISO and PXE installations by using the following methods:

- Kernel arguments: You can use kernel arguments to provide installation-specific information. For example, you can specify the locations of the RHCOS installation files that you uploaded to

your HTTP server and the location of the Ignition config file for the type of node you are installing. For a PXE installation, you can use the **APPEND** parameter to pass the arguments to the kernel of the live installer. For an ISO installation, you can interrupt the live installation boot process to add the kernel arguments. In both installation cases, you can use special **coreos.inst.*** arguments to direct the live installer, as well as standard installation boot arguments for turning standard kernel services on or off.

- Ignition configs: OpenShift Container Platform Ignition config files (***.ign**) are specific to the type of node you are installing. You pass the location of a bootstrap, control plane, or compute node Ignition config file during the RHCOS installation so that it takes effect on first boot. In special cases, you can create a separate, limited Ignition config to pass to the live system. That Ignition config could do a certain set of tasks, such as reporting success to a provisioning system after completing installation. This special Ignition config is consumed by the **coreos-installer** to be applied on first boot of the installed system. Do not provide the standard control plane and compute node Ignition configs to the live ISO directly.
- **coreos-installer**: You can boot the live ISO installer to a shell prompt, which allows you to prepare the permanent system in a variety of ways before first boot. In particular, you can run the **coreos-installer** command to identify various artifacts to include, work with disk partitions, and set up networking. In some cases, you can configure features on the live system and copy them to the installed system.

Whether to use an ISO or PXE install depends on your situation. A PXE install requires an available DHCP service and more preparation, but can make the installation process more automated. An ISO install is a more manual process and can be inconvenient if you are setting up more than a few machines.



NOTE

As of OpenShift Container Platform 4.6, the RHCOS ISO and other installation artifacts provide support for installation on disks with 4K sectors.

15.2.11.1. Installing RHCOS by using an ISO image

You can use an ISO image to install RHCOS on the machines.

Prerequisites

- You have created the Ignition config files for your cluster.
- You have configured suitable network, DNS and load balancing infrastructure.
- You have an HTTP server that can be accessed from your computer, and from the machines that you create.
- You have reviewed the *Advanced RHCOS installation configuration* section for different ways to configure features, such as networking and disk partitioning.

Procedure

1. Obtain the SHA512 digest for each of your Ignition config files. For example, you can use the following on a system running Linux to get the SHA512 digest for your **bootstrap.ign** Ignition config file:

```
$ sha512sum <installation_directory>/bootstrap.ign
```

The digests are provided to the **coreos-installer** in a later step to validate the authenticity of the Ignition config files on the cluster nodes.

2. Upload the bootstrap, control plane, and compute node Ignition config files that the installation program created to your HTTP server. Note the URLs of these files.



IMPORTANT

You can add or change configuration settings in your Ignition configs before saving them to your HTTP server. If you plan to add more compute machines to your cluster after you finish installation, do not delete these files.

3. From the installation host, validate that the Ignition config files are available on the URLs. The following example gets the Ignition config file for the bootstrap node:

```
$ curl -k http://<HTTP_server>/bootstrap.ign 1
```

Example output

```
% Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
           Dload  Upload  Total   Spent    Left  Speed
  0   0   0    0     0    0     0     0  0 --:--:-- --:--:-- --:--:--    0{"ignition":
{"version":"3.2.0"},"passwd":{"users":[{"name":"core","sshAuthorizedKeys":["ssh-rsa...
```

Replace **bootstrap.ign** with **master.ign** or **worker.ign** in the command to validate that the Ignition config files for the control plane and compute nodes are also available.

4. Although it is possible to obtain the RHCOS images that are required for your preferred method of installing operating system instances from the [RHCOS image mirror](#) page, the recommended way to obtain the correct version of your RHCOS images are from the output of **openshift-install** command:

```
$ openshift-install coreos print-stream-json | grep '\.iso[^\.]'
```

Example output

```
"location": "<url>/art/storage/releases/rhcos-4.16-aarch64/<release>/aarch64/rhcos-
<release>-live.aarch64.iso",
"location": "<url>/art/storage/releases/rhcos-4.16-ppc64le/<release>/ppc64le/rhcos-
<release>-live.ppc64le.iso",
"location": "<url>/art/storage/releases/rhcos-4.16-s390x/<release>/s390x/rhcos-<release>-
live.s390x.iso",
"location": "<url>/art/storage/releases/rhcos-4.16/<release>/x86_64/rhcos-<release>-
live.x86_64.iso",
```




IMPORTANT

The RHCOS images might not change with every release of OpenShift Container Platform. You must download images with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Use the image versions that match your OpenShift Container Platform version if they are available. Use only ISO images for this procedure. RHCOS qcow2 images are not supported for this installation type.

ISO file names resemble the following example:

rhcos-<version>-live.<architecture>.iso

5. Use the ISO to start the RHCOS installation. Use one of the following installation options:
 - Burn the ISO image to a disk and boot it directly.
 - Use ISO redirection by using a lights-out management (LOM) interface.
6. Boot the RHCOS ISO image without specifying any options or interrupting the live boot sequence. Wait for the installer to boot into a shell prompt in the RHCOS live environment.



NOTE


It is possible to interrupt the RHCOS installation boot process to add kernel arguments. However, for this ISO procedure you should use the **coreos-installer** command as outlined in the following steps, instead of adding kernel arguments.

7. Run the **coreos-installer** command and specify the options that meet your installation requirements. At a minimum, you must specify the URL that points to the Ignition config file for the node type, and the device that you are installing to:

```
$ sudo coreos-installer install --ignition-url=http://<HTTP_server>/<node_type>.ign <device>
--ignition-hash=sha512-<digest> 1 2
```

- 1 1 You must run the **coreos-installer** command by using **sudo**, because the **core** user does not have the required root privileges to perform the installation.

- 2 The **--ignition-hash** option is required when the Ignition config file is obtained through an HTTP URL to validate the authenticity of the Ignition config file on the cluster node. **<digest>** is the Ignition config file SHA512 digest obtained in a preceding step.




NOTE

If you want to provide your Ignition config files through an HTTPS server that uses TLS, you can add the internal certificate authority (CA) to the system trust store before running **coreos-installer**.

The following example initializes a bootstrap node installation to the **/dev/sda** device. The Ignition config file for the bootstrap node is obtained from an HTTP web server with the IP address 192.168.1.2:

```
$ sudo coreos-installer install --ignition-
url=http://192.168.1.2:80/installation_directory/bootstrap.ign /dev/sda --ignition-hash=sha512-
a5a2d43879223273c9b60af66b44202a1d1248fc01cf156c46d4a79f552b6bad47bc8cc78ddf011
6e80c59d2ea9e32ba53bc807afbca581aa059311def2c3e3b
```

8. Monitor the progress of the RHCOS installation on the console of the machine.



IMPORTANT


Be sure that the installation is successful on each node before commencing with the OpenShift Container Platform installation. Observing the installation process can also help to determine the cause of RHCOS installation issues that might arise.

9. After RHCOS installs, you must reboot the system. During the system reboot, it applies the Ignition config file that you specified.
10. Check the console output to verify that Ignition ran.

Example command

```
Ignition: ran on 2022/03/14 14:48:33 UTC (this boot)
Ignition: user-provided config was applied
```

11. Continue to create the other machines for your cluster.



IMPORTANT

You must create the bootstrap and control plane machines at this time. If the control plane machines are not made schedulable, also create at least two compute machines before you install OpenShift Container Platform.

If the required network, DNS, and load balancer infrastructure are in place, the OpenShift Container Platform bootstrap process begins automatically after the RHCOS nodes have rebooted.



NOTE

RHCOS nodes do not include a default password for the **core** user. You can access the nodes by running **ssh core@<node>.<cluster_name>.<base_domain>** as a user with access to the SSH private key that is paired to the public key that you specified in your **install_config.yaml** file. OpenShift Container Platform 4 cluster nodes running RHCOS are immutable and rely on Operators to apply cluster changes. Accessing cluster nodes by using SSH is not recommended. However, when investigating installation issues, if the OpenShift Container Platform API is not available, or the kubelet is not properly functioning on a target node, SSH access might be required for debugging or disaster recovery.

15.2.11.2. Installing RHCOS by using PXE or iPXE booting

You can use PXE or iPXE booting to install RHCOS on the machines.

Prerequisites

- You have created the Ignition config files for your cluster.
- You have configured suitable network, DNS and load balancing infrastructure.
- You have configured suitable PXE or iPXE infrastructure.

- You have an HTTP server that can be accessed from your computer, and from the machines that you create.
- You have reviewed the *Advanced RHCOS installation configuration* section for different ways to configure features, such as networking and disk partitioning.

Procedure

1. Upload the bootstrap, control plane, and compute node Ignition config files that the installation program created to your HTTP server. Note the URLs of these files.



IMPORTANT

You can add or change configuration settings in your Ignition configs before saving them to your HTTP server. If you plan to add more compute machines to your cluster after you finish installation, do not delete these files.

2. From the installation host, validate that the Ignition config files are available on the URLs. The following example gets the Ignition config file for the bootstrap node:

```
$ curl -k http://<HTTP_server>/bootstrap.ign 1
```

Example output

```
% Total    % Received % Xferd Average Speed   Time    Time     Time  Current
           Dload  Upload   Total   Spent    Left  Speed
  0   0   0   0   0   0   0   0   0  --:--:-- --:--:-- --:--:--    0{"ignition":
{"version":"3.2.0"},"passwd":{"users":[{"name":"core","sshAuthorizedKeys":["ssh-rsa...
```

Replace **bootstrap.ign** with **master.ign** or **worker.ign** in the command to validate that the Ignition config files for the control plane and compute nodes are also available.

3. Although it is possible to obtain the RHCOS **kernel**, **initramfs** and **rootfs** files that are required for your preferred method of installing operating system instances from the [RHCOS image mirror](#) page, the recommended way to obtain the correct version of your RHCOS files are from the output of **openshift-install** command:

```
$ openshift-install coreos print-stream-json | grep -Eo "https.*(kernel-|initramfs-|rootfs-)\w+(\.img)?"
```

Example output

```
"<url>/art/storage/releases/rhcos-4.16-aarch64/<release>/aarch64/rhcos-<release>-live-
kernel-aarch64"
"<url>/art/storage/releases/rhcos-4.16-aarch64/<release>/aarch64/rhcos-<release>-live-
initramfs.aarch64.img"
"<url>/art/storage/releases/rhcos-4.16-aarch64/<release>/aarch64/rhcos-<release>-live-
rootfs.aarch64.img"
"<url>/art/storage/releases/rhcos-4.16-ppc64le/49.84.202110081256-0/ppc64le/rhcos-
<release>-live-kernel-ppc64le"
"<url>/art/storage/releases/rhcos-4.16-ppc64le/<release>/ppc64le/rhcos-<release>-live-
initramfs.ppc64le.img"
"<url>/art/storage/releases/rhcos-4.16-ppc64le/<release>/ppc64le/rhcos-<release>-live-
```

```

rootfs.ppc64le.img"
"<url>/art/storage/releases/rhcos-4.16-s390x/<release>/s390x/rhcos-<release>-live-kernel-
s390x"
"<url>/art/storage/releases/rhcos-4.16-s390x/<release>/s390x/rhcos-<release>-live-
initramfs.s390x.img"
"<url>/art/storage/releases/rhcos-4.16-s390x/<release>/s390x/rhcos-<release>-live-
rootfs.s390x.img"
"<url>/art/storage/releases/rhcos-4.16/<release>/x86_64/rhcos-<release>-live-kernel-
x86_64"
"<url>/art/storage/releases/rhcos-4.16/<release>/x86_64/rhcos-<release>-live-
initramfs.x86_64.img"
"<url>/art/storage/releases/rhcos-4.16/<release>/x86_64/rhcos-<release>-live-
rootfs.x86_64.img"

```



IMPORTANT

The RHCOS artifacts might not change with every release of OpenShift Container Platform. You must download images with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Only use the appropriate **kernel**, **initramfs**, and **rootfs** artifacts described below for this procedure. RHCOS QCOW2 images are not supported for this installation type.

The file names contain the OpenShift Container Platform version number. They resemble the following examples:

- **kernel:** rhcos-<version>-live-kernel-<architecture>
- **initramfs:** rhcos-<version>-live-initramfs.<architecture>.img
- **rootfs:** rhcos-<version>-live-rootfs.<architecture>.img

4. Upload the **rootfs**, **kernel**, and **initramfs** files to your HTTP server.



IMPORTANT

If you plan to add more compute machines to your cluster after you finish installation, do not delete these files.

5. Configure the network boot infrastructure so that the machines boot from their local disks after RHCOS is installed on them.
6. Configure PXE or iPXE installation for the RHCOS images and begin the installation. Modify one of the following example menu entries for your environment and verify that the image and Ignition files are properly accessible:

- For PXE (**x86_64**):

```

DEFAULT pxeboot
TIMEOUT 20
PROMPT 0
LABEL pxeboot
  KERNEL http://<HTTP_server>/rhcos-<version>-live-kernel-<architecture>
  APPEND initrd=http://<HTTP_server>/rhcos-<version>-live-initramfs.

```

1

```
<architecture>.img coreos.live.rootfs_url=http://<HTTP_server>/rhcos-<version>-live-  
rootfs.<architecture>.img coreos.inst.install_dev=/dev/sda  
coreos.inst.ignition_url=http://<HTTP_server>/bootstrap.ign
```

- 1 1 Specify the location of the live **kernel** file that you uploaded to your HTTP server. The URL must be HTTP, TFTP, or FTP; HTTPS and NFS are not supported.
- 2 If you use multiple NICs, specify a single interface in the **ip** option. For example, to use DHCP on a NIC that is named **eno1**, set **ip=eno1:dhcp**.
- 3 Specify the locations of the RHCOS files that you uploaded to your HTTP server. The **initrd** parameter value is the location of the **initramfs** file, the **coreos.live.rootfs_url** parameter value is the location of the **rootfs** file, and the **coreos.inst.ignition_url** parameter value is the location of the bootstrap Ignition config file. You can also add more kernel arguments to the **APPEND** line to configure networking or other boot options.



NOTE

This configuration does not enable serial console access on machines with a graphical console. To configure a different console, add one or more **console=** arguments to the **APPEND** line. For example, add **console=tty0 console=ttyS0** to set the first PC serial port as the primary console and the graphical console as a secondary console. For more information, see [How does one set up a serial terminal and/or console in Red Hat Enterprise Linux?](#) and "Enabling the serial console for PXE and ISO installation" in the "Advanced RHCOS installation configuration" section.

- For iPXE (**x86_64 + aarch64**):

```
kernel http://<HTTP_server>/rhcos-<version>-live-kernel-<architecture> initrd=main  
coreos.live.rootfs_url=http://<HTTP_server>/rhcos-<version>-live-rootfs.  
<architecture>.img coreos.inst.install_dev=/dev/sda  
coreos.inst.ignition_url=http://<HTTP_server>/bootstrap.ign  
initrd --name main http://<HTTP_server>/rhcos-<version>-live-initramfs.  
<architecture>.img  
boot
```

- 1 Specify the locations of the RHCOS files that you uploaded to your HTTP server. The **kernel** parameter value is the location of the **kernel** file, the **initrd=main** argument is needed for booting on UEFI systems, the **coreos.live.rootfs_url** parameter value is the location of the **rootfs** file, and the **coreos.inst.ignition_url** parameter value is the location of the bootstrap Ignition config file.
- 2 If you use multiple NICs, specify a single interface in the **ip** option. For example, to use DHCP on a NIC that is named **eno1**, set **ip=eno1:dhcp**.
- 3 Specify the location of the **initramfs** file that you uploaded to your HTTP server.

**NOTE**

This configuration does not enable serial console access on machines with a graphical console. To configure a different console, add one or more **console=** arguments to the **kernel** line. For example, add **console=tty0 console=ttyS0** to set the first PC serial port as the primary console and the graphical console as a secondary console. For more information, see [How does one set up a serial terminal and/or console in Red Hat Enterprise Linux?](#) and "Enabling the serial console for PXE and ISO installation" in the "Advanced RHCOS installation configuration" section.

**NOTE**

To network boot the CoreOS **kernel** on **aarch64** architecture, you need to use a version of iPXE build with the **IMAGE_GZIP** option enabled. See [IMAGE_GZIP option in iPXE](#).

- For PXE (with UEFI and Grub as second stage) on **aarch64**:


```

menuentry 'Install CoreOS' {
  linux rhcos-<version>-live-kernel-<architecture>
  coreos.live.rootfs_url=http://<HTTP_server>/rhcos-<version>-live-rootfs.
  <architecture>.img coreos.inst.install_dev=/dev/sda
  coreos.inst.ignition_url=http://<HTTP_server>/bootstrap.ign 1 2
  initrd rhcos-<version>-live-initramfs.<architecture>.img 3
}

```

- 1** Specify the locations of the RHCOS files that you uploaded to your HTTP/TFTP server. The **kernel** parameter value is the location of the **kernel** file on your TFTP server. The **coreos.live.rootfs_url** parameter value is the location of the **rootfs** file, and the **coreos.inst.ignition_url** parameter value is the location of the bootstrap Ignition config file on your HTTP Server.
- 2** If you use multiple NICs, specify a single interface in the **ip** option. For example, to use DHCP on a NIC that is named **eno1**, set **ip=eno1:dhcp**.
- 3** Specify the location of the **initramfs** file that you uploaded to your TFTP server.

7. Monitor the progress of the RHCOS installation on the console of the machine.

**IMPORTANT**

Be sure that the installation is successful on each node before commencing with the OpenShift Container Platform installation. Observing the installation process can also help to determine the cause of RHCOS installation issues that might arise.


8. After RHCOS installs, the system reboots. During reboot, the system applies the Ignition config file that you specified.
9. Check the console output to verify that Ignition ran.

Example command

-

Ignition: ran on 2022/03/14 14:48:33 UTC (this boot)
 Ignition: user-provided config was applied

10. Continue to create the machines for your cluster.



IMPORTANT

You must create the bootstrap and control plane machines at this time. If the control plane machines are not made schedulable, also create at least two compute machines before you install the cluster.

If the required network, DNS, and load balancer infrastructure are in place, the OpenShift Container Platform bootstrap process begins automatically after the RHCOS nodes have rebooted.



NOTE

RHCOS nodes do not include a default password for the **core** user. You can access the nodes by running **ssh core@<node>.<cluster_name>.<base_domain>** as a user with access to the SSH private key that is paired to the public key that you specified in your **install_config.yaml** file. OpenShift Container Platform 4 cluster nodes running RHCOS are immutable and rely on Operators to apply cluster changes. Accessing cluster nodes by using SSH is not recommended. However, when investigating installation issues, if the OpenShift Container Platform API is not available, or the kubelet is not properly functioning on a target node, SSH access might be required for debugging or disaster recovery.

15.2.11.3. Advanced RHCOS installation configuration

A key benefit for manually provisioning the Red Hat Enterprise Linux CoreOS (RHCOS) nodes for OpenShift Container Platform is to be able to do configuration that is not available through default OpenShift Container Platform installation methods. This section describes some of the configurations that you can do using techniques that include:

- Passing kernel arguments to the live installer
- Running **coreos-installer** manually from the live system
- Customizing a live ISO or PXE boot image

The advanced configuration topics for manual Red Hat Enterprise Linux CoreOS (RHCOS) installations detailed in this section relate to disk partitioning, networking, and using Ignition configs in different ways.

15.2.11.3.1. Using advanced networking options for PXE and ISO installations

Networking for OpenShift Container Platform nodes uses DHCP by default to gather all necessary configuration settings. To set up static IP addresses or configure special settings, such as bonding, you can do one of the following:

- Pass special kernel parameters when you boot the live installer.
- Use a machine config to copy networking files to the installed system.

- Configure networking from a live installer shell prompt, then copy those settings to the installed system so that they take effect when the installed system first boots.

To configure a PXE or iPXE installation, use one of the following options:


- See the "Advanced RHCOS installation reference" tables.
- Use a machine config to copy networking files to the installed system.

To configure an ISO installation, use the following procedure.

Procedure

1. Boot the ISO installer.
2. From the live system shell prompt, configure networking for the live system using available RHEL tools, such as **nmcli** or **nmtui**.
3. Run the **coreos-installer** command to install the system, adding the **--copy-network** option to copy networking configuration. For example:

```
$ sudo coreos-installer install --copy-network \  
--ignition-url=http://host/worker.ign /dev/disk/by-id/scsi-<serial_number>
```



IMPORTANT

The **--copy-network** option only copies networking configuration found under **/etc/NetworkManager/system-connections**. In particular, it does not copy the system hostname.


4. Reboot into the installed system.

Additional resources

- See [Getting started with nmcli](#) and [Getting started with nmtui](#) in the RHEL 8 documentation for more information about the **nmcli** and **nmtui** tools.

15.2.11.3.2. Disk partitioning

Disk partitions are created on OpenShift Container Platform cluster nodes during the Red Hat Enterprise Linux CoreOS (RHCOS) installation. Each RHCOS node of a particular architecture uses the same partition layout, unless you override the default partitioning configuration. During the RHCOS installation, the size of the root file system is increased to use any remaining available space on the target device.



IMPORTANT

The use of a custom partition scheme on your node might result in OpenShift Container Platform not monitoring or alerting on some node partitions. If you override the default partitioning, see [Understanding OpenShift File System Monitoring \(eviction conditions\)](#) for more information about how OpenShift Container Platform monitors your host file systems.

OpenShift Container Platform monitors the following two filesystem identifiers:

- **nodefs**, which is the filesystem that contains **/var/lib/kubelet**
- **imagefs**, which is the filesystem that contains **/var/lib/containers**

For the default partition scheme, **nodefs** and **imagefs** monitor the same root filesystem, **/**.

To override the default partitioning when installing RHCOS on an OpenShift Container Platform cluster node, you must create separate partitions. Consider a situation where you want to add a separate storage partition for your containers and container images. For example, by mounting **/var/lib/containers** in a separate partition, the kubelet separately monitors **/var/lib/containers** as the **imagefs** directory and the root file system as the **nodefs** directory.



IMPORTANT


If you have resized your disk size to host a larger file system, consider creating a separate **/var/lib/containers** partition. Consider resizing a disk that has an **xfs** format to reduce CPU time issues caused by a high number of allocation groups.

15.2.11.3.2.1. Creating a separate **/var** partition

In general, you should use the default disk partitioning that is created during the RHCOS installation. However, there are cases where you might want to create a separate partition for a directory that you expect to grow.

OpenShift Container Platform supports the addition of a single partition to attach storage to either the **/var** directory or a subdirectory of **/var**. For example:

- **/var/lib/containers**: Holds container-related content that can grow as more images and containers are added to a system.
- **/var/lib/etcd**: Holds data that you might want to keep separate for purposes such as performance optimization of etcd storage.
- **/var**: Holds data that you might want to keep separate for purposes such as auditing.



IMPORTANT

For disk sizes larger than 100GB, and especially larger than 1TB, create a separate **/var** partition.

Storing the contents of a **/var** directory separately makes it easier to grow storage for those areas as needed and reinstall OpenShift Container Platform at a later date and keep that data intact. With this method, you will not have to pull all your containers again, nor will you have to copy massive log files when you update systems.

The use of a separate partition for the **/var** directory or a subdirectory of **/var** also prevents data growth in the partitioned directory from filling up the root file system.

The following procedure sets up a separate **/var** partition by adding a machine config manifest that is wrapped into the Ignition config file for a node type during the preparation phase of an installation.

Procedure

1. On your installation host, change to the directory that contains the OpenShift Container Platform installation program and generate the Kubernetes manifests for the cluster:

```
$ openshift-install create manifests --dir <installation_directory>
```

2. Create a Butane config that configures the additional partition. For example, name the file **\$HOME/clusterconfig/98-var-partition.bu**, change the disk device name to the name of the storage device on the **worker** systems, and set the storage size as appropriate. This example places the **/var** directory on a separate partition:

```
variant: openshift
version: 4.16.0
metadata:
  labels:
    machineconfiguration.openshift.io/role: worker
  name: 98-var-partition
storage:
  disks:
    - device: /dev/disk/by-id/<device_name> 1
    partitions:
      - label: var
        start_mib: <partition_start_offset> 2
        size_mib: <partition_size> 3
        number: 5
    filesystems:
      - device: /dev/disk/by-partlabel/var
        path: /var
        format: xfs
        mount_options: [defaults, prjquota] 4
        with_mount_unit: true
```

- 1 The storage device name of the disk that you want to partition.
- 2 When adding a data partition to the boot disk, a minimum offset value of 25000 mebibytes is recommended. The root file system is automatically resized to fill all available space up to the specified offset. If no offset value is specified, or if the specified value is smaller than the recommended minimum, the resulting root file system will be too small, and future reinstalls of RHCOS might overwrite the beginning of the data partition.
- 3 The size of the data partition in mebibytes.
- 4 The **prjquota** mount option must be enabled for filesystems used for container storage.



NOTE

When creating a separate **/var** partition, you cannot use different instance types for compute nodes, if the different instance types do not have the same device name.

3. Create a manifest from the Butane config and save it to the **clusterconfig/openshift** directory. For example, run the following command:

```
$ butane $HOME/clusterconfig/98-var-partition.bu -o $HOME/clusterconfig/openshift/98-var-partition.yaml
```

4. Create the Ignition config files:


```
$ openshift-install create ignition-configs --dir <installation_directory> 1
```

1 For **<installation_directory>**, specify the same installation directory.

Ignition config files are created for the bootstrap, control plane, and compute nodes in the installation directory:

```
.
├── auth
│   ├── kubeadmin-password
│   └── kubeconfig
├── bootstrap.ign
├── master.ign
├── metadata.json
└── worker.ign
```

The files in the **<installation_directory>/manifest** and **<installation_directory>/openshift** directories are wrapped into the Ignition config files, including the file that contains the **98-var-partition** custom **MachineConfig** object.

Next steps

- You can apply the custom disk partitioning by referencing the Ignition config files during the RHCOS installations.

15.2.11.3.2.2. Retaining existing partitions

For an ISO installation, you can add options to the **coreos-installer** command that cause the installer to maintain one or more existing partitions. For a PXE installation, you can add **coreos.inst.*** options to the **APPEND** parameter to preserve partitions.

Saved partitions might be data partitions from an existing OpenShift Container Platform system. You can identify the disk partitions you want to keep either by partition label or by number.



NOTE

If you save existing partitions, and those partitions do not leave enough space for RHCOS, the installation will fail without damaging the saved partitions.

Retaining existing partitions during an ISO installation

This example preserves any partition in which the partition label begins with **data** (**data***):

```
# coreos-installer install --ignition-url http://10.0.2.2:8080/user.ign \
  --save-partlabel 'data*' /dev/disk/by-id/scsi-<serial_number>
```

The following example illustrates running the **coreos-installer** in a way that preserves the sixth (6) partition on the disk:

```
# coreos-installer install --ignition-url http://10.0.2.2:8080/user.ign \
  --save-partindex 6 /dev/disk/by-id/scsi-<serial_number>
```

This example preserves partitions 5 and higher:

```
# coreos-installer install --ignition-url http://10.0.2.2:8080/user.ign
--save-partindex 5- /dev/disk/by-id/scsi-<serial_number>
```

In the previous examples where partition saving is used, **coreos-installer** recreates the partition immediately.

Retaining existing partitions during a PXE installation

This **APPEND** option preserves any partition in which the partition label begins with 'data' ('data*'):

```
coreos.inst.save_partlabel=data*
```

This **APPEND** option preserves partitions 5 and higher:

```
coreos.inst.save_partindex=5-
```

This **APPEND** option preserves partition 6:

```
coreos.inst.save_partindex=6
```

15.2.11.3.3. Identifying Ignition configs

When doing an RHCOS manual installation, there are two types of Ignition configs that you can provide, with different reasons for providing each one:

- **Permanent install Ignition config:** Every manual RHCOS installation needs to pass one of the Ignition config files generated by **openshift-installer**, such as **bootstrap.ign**, **master.ign** and **worker.ign**, to carry out the installation.



IMPORTANT

It is not recommended to modify these Ignition config files directly. You can update the manifest files that are wrapped into the Ignition config files, as outlined in examples in the preceding sections.

For PXE installations, you pass the Ignition configs on the **APPEND** line using the **coreos.inst.ignition_url=** option. For ISO installations, after the ISO boots to the shell prompt, you identify the Ignition config on the **coreos-installer** command line with the **--ignition-url=** option. In both cases, only HTTP and HTTPS protocols are supported.

- **Live install Ignition config:** This type can be created by using the **coreos-installer customize** subcommand and its various options. With this method, the Ignition config passes to the live install medium, runs immediately upon booting, and performs setup tasks before or after the RHCOS system installs to disk. This method should only be used for performing tasks that must be done once and not applied again later, such as with advanced partitioning that cannot be done using a machine config.

For PXE or ISO boots, you can create the Ignition config and **APPEND** the **ignition.config.url=** option to identify the location of the Ignition config. You also need to append **ignition.firstboot** **ignition.platform.id=metal** or the **ignition.config.url** option will be ignored.

15.2.11.3.4. Default console configuration

Red Hat Enterprise Linux CoreOS (RHCOS) nodes installed from an OpenShift Container Platform 4.16

boot image use a default console that is meant to accommodate most virtualized and bare metal setups. Different cloud and virtualization platforms may use different default settings depending on the chosen architecture. Bare metal installations use the kernel default settings which typically means the graphical console is the primary console and the serial console is disabled.

The default consoles may not match your specific hardware configuration or you might have specific needs that require you to adjust the default console. For example:

- You want to access the emergency shell on the console for debugging purposes.
- Your cloud platform does not provide interactive access to the graphical console, but provides a serial console.
- You want to enable multiple consoles.

Console configuration is inherited from the boot image. This means that new nodes in existing clusters are unaffected by changes to the default console.

You can configure the console for bare metal installations in the following ways:

- Using **coreos-installer** manually on the command line.
- Using the **coreos-installer iso customize** or **coreos-installer pxe customize** subcommands with the **--dest-console** option to create a custom image that automates the process.



NOTE

For advanced customization, perform console configuration using the **coreos-installer iso** or **coreos-installer pxe** subcommands, and not kernel arguments.

15.2.11.3.5. Enabling the serial console for PXE and ISO installations

By default, the Red Hat Enterprise Linux CoreOS (RHCOS) serial console is disabled and all output is written to the graphical console. You can enable the serial console for an ISO installation and reconfigure the bootloader so that output is sent to both the serial console and the graphical console.

Procedure


1. Boot the ISO installer.
2. Run the **coreos-installer** command to install the system, adding the **--console** option once to specify the graphical console, and a second time to specify the serial console:

```
$ coreos-installer install \
  --console=tty0 \1
  --console=ttyS0,<options> \2
  --ignition-url=http://host/worker.ign /dev/disk/by-id/scsi-<serial_number>
```

1 The desired secondary console. In this case, the graphical console. Omitting this option will disable the graphical console.

2 The desired primary console. In this case the serial console. The **options** field defines the baud rate and other settings. A common value for this field is **11520n8**. If no options are provided, the default kernel value of **9600n8** is used. For more information on the format of this option, see [Linux kernel serial console](#) documentation.

3. Reboot into the installed system.



NOTE

A similar outcome can be obtained by using the **coreos-installer install --append-karg** option, and specifying the console with **console=**. However, this will only set the console for the kernel and not the bootloader.

To configure a PXE installation, make sure the **coreos.inst.install_dev** kernel command line option is omitted, and use the shell prompt to run **coreos-installer** manually using the above ISO installation procedure.

15.2.11.3.6. Customizing a live RHCOS ISO or PXE install

You can use the live ISO image or PXE environment to install RHCOS by injecting an Ignition config file directly into the image. This creates a customized image that you can use to provision your system.

For an ISO image, the mechanism to do this is the **coreos-installer iso customize** subcommand, which modifies the **.iso** file with your configuration. Similarly, the mechanism for a PXE environment is the **coreos-installer pxe customize** subcommand, which creates a new **initramfs** file that includes your customizations.

The **customize** subcommand is a general purpose tool that can embed other types of customizations as well. The following tasks are examples of some of the more common customizations:

- Inject custom CA certificates for when corporate security policy requires their use.
- Configure network settings without the need for kernel arguments.
- Embed arbitrary preinstall and post-install scripts or binaries.

15.2.11.3.7. Customizing a live RHCOS ISO image

You can customize a live RHCOS ISO image directly with the **coreos-installer iso customize** subcommand. When you boot the ISO image, the customizations are applied automatically.

You can use this feature to configure the ISO image to automatically install RHCOS.

Procedure

1. Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.
2. Retrieve the RHCOS ISO image from the [RHCOS image mirror](#) page and the Ignition config file, and then run the following command to inject the Ignition config directly into the ISO image:

```
$ coreos-installer iso customize rhcos-<version>-live.x86_64.iso \
  --dest-ignition bootstrap.ign \ 1
  --dest-device /dev/disk/by-id/scsi-<serial_number> 2
```

- 1 The Ignition config file that is generated from the **openshift-installer** installation program.
- 2 When you specify this option, the ISO image automatically runs an installation. Otherwise, the image remains configured for installation, but does not install automatically unless you specify the **coreos.inst.install_dev** kernel argument.

- Optional: To remove the ISO image customizations and return the image to its pristine state, run:

```
$ coreos-installer iso reset rhcos-<version>-live.x86_64.iso
```

You can now re-customize the live ISO image or use it in its pristine state.

Applying your customizations affects every subsequent boot of RHCOS.

15.2.11.3.7.1. Modifying a live install ISO image to enable the serial console

On clusters installed with OpenShift Container Platform 4.12 and above, the serial console is disabled by default and all output is written to the graphical console. You can enable the serial console with the following procedure.

Procedure

- Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.
- Retrieve the RHCOS ISO image from the [RHCOS image mirror](#) page and run the following command to customize the ISO image to enable the serial console to receive output:

```
$ coreos-installer iso customize rhcos-<version>-live.x86_64.iso \
  --dest-ignition <path> \ 1
  --dest-console tty0 \ 2
  --dest-console ttyS0,<options> \ 3
  --dest-device /dev/disk/by-id/scsi-<serial_number> 4
```

- The location of the Ignition config to install.
- The desired secondary console. In this case, the graphical console. Omitting this option will disable the graphical console.
- The desired primary console. In this case, the serial console. The **options** field defines the baud rate and other settings. A common value for this field is **115200n8**. If no options are provided, the default kernel value of **9600n8** is used. For more information on the format of this option, see the [Linux kernel serial console](#) documentation.
- The specified disk to install to. If you omit this option, the ISO image automatically runs the installation program which will fail unless you also specify the **coreos.inst.install_dev** kernel argument.



NOTE

The **--dest-console** option affects the installed system and not the live ISO system. To modify the console for a live ISO system, use the **--live-karg-append** option and specify the console with **console=**.

Your customizations are applied and affect every subsequent boot of the ISO image.

- Optional: To remove the ISO image customizations and return the image to its original state, run the following command:

```
$ coreos-installer iso reset rhcos-<version>-live.x86_64.iso
```

You can now recustomize the live ISO image or use it in its original state.

15.2.11.3.7.2. Modifying a live install ISO image to use a custom certificate authority

You can provide certificate authority (CA) certificates to Ignition with the **--ignition-ca** flag of the **customize** subcommand. You can use the CA certificates during both the installation boot and when provisioning the installed system.



NOTE

Custom CA certificates affect how Ignition fetches remote resources but they do not affect the certificates installed onto the system.

Procedure

1. Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.
2. Retrieve the RHCOS ISO image from the [RHCOS image mirror](#) page and run the following command to customize the ISO image for use with a custom CA:

```
$ coreos-installer iso customize rhcos-<version>-live.x86_64.iso --ignition-ca cert.pem
```



IMPORTANT

The **coreos.inst.ignition_url** kernel parameter does not work with the **--ignition-ca** flag. You must use the **--dest-ignition** flag to create a customized image for each cluster.

Applying your custom CA certificate affects every subsequent boot of RHCOS.

15.2.11.3.7.3. Modifying a live install ISO image with customized network settings

You can embed a NetworkManager keyfile into the live ISO image and pass it through to the installed system with the **--network-keyfile** flag of the **customize** subcommand.



WARNING

When creating a connection profile, you must use a **.nmconnection** filename extension in the filename of the connection profile. If you do not use a **.nmconnection** filename extension, the cluster will apply the connection profile to the live environment, but it will not apply the configuration when the cluster first boots up the nodes, resulting in a setup that does not work.

Procedure

1. Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.

2. Create a connection profile for a bonded interface. For example, create the **bond0.nmconnection** file in your local directory with the following content:

```
[connection]
id=bond0
type=bond
interface-name=bond0
multi-connect=1
permissions=

[ethernet]
mac-address-blacklist=

[bond]
miimon=100
mode=active-backup

[ipv4]
method=auto

[ipv6]
method=auto

[proxy]
```

3. Create a connection profile for a secondary interface to add to the bond. For example, create the **bond0-proxy-em1.nmconnection** file in your local directory with the following content:

```
[connection]
id=em1
type=ethernet
interface-name=em1
master=bond0
multi-connect=1
permissions=
slave-type=bond

[ethernet]
mac-address-blacklist=
```

4. Create a connection profile for a secondary interface to add to the bond. For example, create the **bond0-proxy-em2.nmconnection** file in your local directory with the following content:

```
[connection]
id=em2
type=ethernet
interface-name=em2
master=bond0
multi-connect=1
permissions=
slave-type=bond

[ethernet]
mac-address-blacklist=
```

- Retrieve the RHCOS ISO image from the [RHCOS image mirror](#) page and run the following command to customize the ISO image with your configured networking:

```
$ coreos-installer iso customize rhcos-<version>-live.x86_64.iso \
  --network-keyfile bond0.nmconnection \
  --network-keyfile bond0-proxy-em1.nmconnection \
  --network-keyfile bond0-proxy-em2.nmconnection
```

Network settings are applied to the live system and are carried over to the destination system.

15.2.11.3.7.4. Customizing a live install ISO image for an iSCSI boot device

You can set the iSCSI target and initiator values for automatic mounting, booting and configuration using a customized version of the live RHCOS image.

Prerequisites

- You have an iSCSI target you want to install RHCOS on.

Procedure

- Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.
- Retrieve the RHCOS ISO image from the [RHCOS image mirror](#) page and run the following command to customize the ISO image with the following information:

```
$ coreos-installer iso customize \
  --pre-install mount-iscsi.sh \ 1
  --post-install unmount-iscsi.sh \ 2
  --dest-device /dev/disk/by-path/<IP_address>:<port>-iscsi-<target_iqn>-lun-<lun> \ 3
  --dest-ignition config.ign \ 4
  --dest-karg-append rd.iscsi.initiator=<initiator_iqn> \ 5
  --dest-karg-append netroot=<target_iqn> \ 6
  -o custom.iso rhcos-<version>-live.x86_64.iso
```

- The script that gets run before installation. It should contain the **iscsiadm** commands for mounting the iSCSI target and any commands enabling multipathing.
- The script that gets run after installation. It should contain the command **iscsiadm --mode node --logout=all**.
- The location of the destination system. You must provide the IP address of the target portal, the associated port number, the target iSCSI node in IQN format, and the iSCSI logical unit number (LUN).
- The Ignition configuration for the destination system.
- The iSCSI initiator, or client, name in IQN format. The initiator forms a session to connect to the iSCSI target.
- The the iSCSI target, or server, name in IQN format.

For more information about the iSCSI options supported by **dracut**, see the [dracut.cmdline manual page](#).

15.2.11.3.7.5. Customizing a live install ISO image for an iSCSI boot device with iBFT

You can set the iSCSI target and initiator values for automatic mounting, booting and configuration using a customized version of the live RHCOS image.

Prerequisites

1. You have an iSCSI target you want to install RHCOS on.
2. Optional: you have multipathed your iSCSI target.

Procedure

1. Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.
2. Retrieve the RHCOS ISO image from the [RHCOS image mirror](#) page and run the following command to customize the ISO image with the following information:

```
$ coreos-installer iso customize \
  --pre-install mount-iscsi.sh \ 1
  --post-install unmount-iscsi.sh \ 2
  --dest-device /dev/mapper/mpatha \ 3
  --dest-ignition config.ign \ 4
  --dest-karg-append rd.iscsi.firmware=1 \ 5
  --dest-karg-append rd.multipath=default \ 6
  -o custom.iso rhcos-<version>-live.x86_64.iso
```

- 1 The script that gets run before installation. It should contain the **iscsiadm** commands for mounting the iSCSI target and any commands enabling multipathing.
- 2 The script that gets run after installation. It should contain the command **iscsiadm --mode node --logout=all**.
- 3 The path to the device. If you are using multipath, the multipath device, **/dev/mapper/mpatha**. If there are multiple multipath devices connected, or to be explicit, you can use the World Wide Name (WWN) symlink available in **/dev/disk/by-path**.
- 4 The Ignition configuration for the destination system.
- 5 The iSCSI parameter is read from the BIOS firmware.
- 6 Optional: include this parameter if you are enabling multipathing.

For more information about the iSCSI options supported by **dracut**, see the [dracut.cmdline manual page](#).

15.2.11.3.8. Customizing a live RHCOS PXE environment

You can customize a live RHCOS PXE environment directly with the **coreos-installer pxe customize** subcommand. When you boot the PXE environment, the customizations are applied automatically.

You can use this feature to configure the PXE environment to automatically install RHCOS.

Procedure

1. Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.
2. Retrieve the RHCOS **kernel**, **initramfs** and **rootfs** files from the [RHCOS image mirror](#) page and the Ignition config file, and then run the following command to create a new **initramfs** file that contains the customizations from your Ignition config:

```
$ coreos-installer pxe customize rhcos-<version>-live-initramfs.x86_64.img \
  --dest-ignition bootstrap.ign \ 1
  --dest-device /dev/disk/by-id/scsi-<serial_number> \ 2
  -o rhcos-<version>-custom-initramfs.x86_64.img 3
```

- 1 The Ignition config file that is generated from **openshift-installer**.
- 2 When you specify this option, the PXE environment automatically runs an install. Otherwise, the image remains configured for installing, but does not do so automatically unless you specify the **coreos.inst.install_dev** kernel argument.
- 3 Use the customized **initramfs** file in your PXE configuration. Add the **ignition.firstboot** and **ignition.platform.id=metal** kernel arguments if they are not already present.

Applying your customizations affects every subsequent boot of RHCOS.

15.2.11.3.8.1. Modifying a live install PXE environment to enable the serial console

On clusters installed with OpenShift Container Platform 4.12 and above, the serial console is disabled by default and all output is written to the graphical console. You can enable the serial console with the following procedure.

Procedure

1. Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.
2. Retrieve the RHCOS **kernel**, **initramfs** and **rootfs** files from the [RHCOS image mirror](#) page and the Ignition config file, and then run the following command to create a new customized **initramfs** file that enables the serial console to receive output:

```
$ coreos-installer pxe customize rhcos-<version>-live-initramfs.x86_64.img \
  --dest-ignition <path> \ 1
  --dest-console tty0 \ 2
  --dest-console ttyS0,<options> \ 3
  --dest-device /dev/disk/by-id/scsi-<serial_number> \ 4
  -o rhcos-<version>-custom-initramfs.x86_64.img 5
```

- 1 The location of the Ignition config to install.
- 2 The desired secondary console. In this case, the graphical console. Omitting this option will disable the graphical console.
- 3 The desired primary console. In this case, the serial console. The **options** field defines the baud rate and other settings. A common value for this field is **115200n8**. If no options are provided, the default kernel value of **9600n8** is used. For more information on the format of this option, see the [Linux kernel serial console](#) documentation.

- 4 The specified disk to install to. If you omit this option, the PXE environment automatically runs the installer which will fail unless you also specify the **coreos.inst.install_dev** kernel argument.
- 5 Use the customized **initramfs** file in your PXE configuration. Add the **ignition.firstboot** and **ignition.platform.id=metal** kernel arguments if they are not already present.

Your customizations are applied and affect every subsequent boot of the PXE environment.

15.2.11.3.8.2. Modifying a live install PXE environment to use a custom certificate authority

You can provide certificate authority (CA) certificates to Ignition with the **--ignition-ca** flag of the **customize** subcommand. You can use the CA certificates during both the installation boot and when provisioning the installed system.



NOTE

Custom CA certificates affect how Ignition fetches remote resources but they do not affect the certificates installed onto the system.

Procedure

1. Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.
2. Retrieve the RHCOS **kernel**, **initramfs** and **rootfs** files from the [RHCOS image mirror](#) page and run the following command to create a new customized **initramfs** file for use with a custom CA:

```
$ coreos-installer pxe customize rhcos-<version>-live-initramfs.x86_64.img \
--ignition-ca cert.pem \
-o rhcos-<version>-custom-initramfs.x86_64.img
```

3. Use the customized **initramfs** file in your PXE configuration. Add the **ignition.firstboot** and **ignition.platform.id=metal** kernel arguments if they are not already present.



IMPORTANT

The **coreos.inst.ignition_url** kernel parameter does not work with the **--ignition-ca** flag. You must use the **--dest-ignition** flag to create a customized image for each cluster.

Applying your custom CA certificate affects every subsequent boot of RHCOS.

15.2.11.3.8.3. Modifying a live install PXE environment with customized network settings

You can embed a NetworkManager keyfile into the live PXE environment and pass it through to the installed system with the **--network-keyfile** flag of the **customize** subcommand.

**WARNING**

When creating a connection profile, you must use a **.nmconnection** filename extension in the filename of the connection profile. If you do not use a **.nmconnection** filename extension, the cluster will apply the connection profile to the live environment, but it will not apply the configuration when the cluster first boots up the nodes, resulting in a setup that does not work.

Procedure

1. Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.
2. Create a connection profile for a bonded interface. For example, create the **bond0.nmconnection** file in your local directory with the following content:

```
[connection]
id=bond0
type=bond
interface-name=bond0
multi-connect=1
permissions=

[ethernet]
mac-address-blacklist=

[bond]
miimon=100
mode=active-backup

[ipv4]
method=auto

[ipv6]
method=auto

[proxy]
```

3. Create a connection profile for a secondary interface to add to the bond. For example, create the **bond0-proxy-em1.nmconnection** file in your local directory with the following content:

```
[connection]
id=em1
type=ethernet
interface-name=em1
master=bond0
multi-connect=1
permissions=
slave-type=bond

[ethernet]
mac-address-blacklist=
```

4. Create a connection profile for a secondary interface to add to the bond. For example, create the **bond0-proxy-em2.nmconnection** file in your local directory with the following content:

```
[connection]
id=em2
type=ethernet
interface-name=em2
master=bond0
multi-connect=1
permissions=
slave-type=bond

[ethernet]
mac-address-blacklist=
```

5. Retrieve the RHCOS **kernel**, **initramfs** and **rootfs** files from the [RHCOS image mirror](#) page and run the following command to create a new customized **initramfs** file that contains your configured networking:

```
$ coreos-installer pxe customize rhcos-<version>-live-initramfs.x86_64.img \
--network-keyfile bond0.nmconnection \
--network-keyfile bond0-proxy-em1.nmconnection \
--network-keyfile bond0-proxy-em2.nmconnection \
-o rhcos-<version>-custom-initramfs.x86_64.img
```

6. Use the customized **initramfs** file in your PXE configuration. Add the **ignition.firstboot** and **ignition.platform.id=metal** kernel arguments if they are not already present. Network settings are applied to the live system and are carried over to the destination system.

15.2.11.3.8.4. Customizing a live install PXE environment for an iSCSI boot device

You can set the iSCSI target and initiator values for automatic mounting, booting and configuration using a customized version of the live RHCOS image.

Prerequisites

1. You have an iSCSI target you want to install RHCOS on.

Procedure

1. Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.
2. Retrieve the RHCOS **kernel**, **initramfs** and **rootfs** files from the [RHCOS image mirror](#) page and run the following command to create a new customized **initramfs** file with the following information:

```
$ coreos-installer pxe customize \
--pre-install mount-iscsi.sh \ 1
--post-install unmount-iscsi.sh \ 2
--dest-device /dev/disk/by-path/<IP_address>:<port>-iscsi-<target_iqn>-lun-<lun> \ 3
--dest-ignition config.ign \ 4
--dest-karg-append rd.iscsi.initiator=<initiator_iqn> \ 5
--dest-karg-append netroot=<target_iqn> \ 6
-o custom.img rhcos-<version>-live-initramfs.x86_64.img
```

-
- 1 The script that gets run before installation. It should contain the **iscsiadm** commands for mounting the iSCSI target and any commands enabling multipathing.
- 2 The script that gets run after installation. It should contain the command **iscsiadm --mode node --logout=all**.
- 3 The location of the destination system. You must provide the IP address of the target portal, the associated port number, the target iSCSI node in IQN format, and the iSCSI logical unit number (LUN).
- 4 The Ignition configuration for the destination system.
- 5 The iSCSI initiator, or client, name in IQN format. The initiator forms a session to connect to the iSCSI target.
- 6 The the iSCSI target, or server, name in IQN format.

For more information about the iSCSI options supported by **dracut**, see the [dracut.cmdline manual page](#).

15.2.11.3.8.5. Customizing a live install PXE environment for an iSCSI boot device with iBFT

You can set the iSCSI target and initiator values for automatic mounting, booting and configuration using a customized version of the live RHCOS image.

Prerequisites

1. You have an iSCSI target you want to install RHCOS on.
2. Optional: you have multipathed your iSCSI target.

Procedure

1. Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.
2. Retrieve the RHCOS **kernel**, **initramfs** and **rootfs** files from the [RHCOS image mirror](#) page and run the following command to create a new customized **initramfs** file with the following information:

```
$ coreos-installer pxe customize \
  --pre-install mount-iscsi.sh \ 1
  --post-install unmount-iscsi.sh \ 2
  --dest-device /dev/mapper/mpatha \ 3
  --dest-ignition config.ign \ 4
  --dest-karg-append rd.iscsi.firmware=1 \ 5
  --dest-karg-append rd.multipath=default \ 6
  -o custom.img rhcos-<version>-live-initramfs.x86_64.img
```

- 1 The script that gets run before installation. It should contain the **iscsiadm** commands for mounting the iSCSI target.
- 2 The script that gets run after installation. It should contain the command **iscsiadm --mode node --logout=all**.

- 3 The path to the device. If you are using multipath, the multipath device, **/dev/mapper/mpatha**, If there are multiple multipath devices connected, or to be explicit,
- 4 The Ignition configuration for the destination system.
- 5 The iSCSI parameter is read from the BIOS firmware.
- 6 Optional: include this parameter if you are enabling multipathing.


For more information about the iSCSI options supported by **dracut**, see the [dracut.cmdline manual page](#).

15.2.11.3.9. Advanced RHCOS installation reference

This section illustrates the networking configuration and other advanced options that allow you to modify the Red Hat Enterprise Linux CoreOS (RHCOS) manual installation process. The following tables describe the kernel arguments and command-line options you can use with the RHCOS live installer and the **coreos-installer** command.

15.2.11.3.9.1. Networking and bonding options for ISO installations

If you install RHCOS from an ISO image, you can add kernel arguments manually when you boot the image to configure networking for a node. If no networking arguments are specified, DHCP is activated in the initramfs when RHCOS detects that networking is required to fetch the Ignition config file.



IMPORTANT

When adding networking arguments manually, you must also add the **rd.neednet=1** kernel argument to bring the network up in the initramfs.

The following information provides examples for configuring networking and bonding on your RHCOS nodes for ISO installations. The examples describe how to use the **ip=**, **nameserver=**, and **bond=** kernel arguments.



NOTE

Ordering is important when adding the kernel arguments: **ip=**, **nameserver=**, and then **bond=**.

The networking options are passed to the **dracut** tool during system boot. For more information about the networking options supported by **dracut**, see the [dracut.cmdline manual page](#).

The following examples are the networking options for ISO installation.

Configuring DHCP or static IP addresses

To configure an IP address, either use DHCP (**ip=dhcp**) or set an individual static IP address (**ip=<host_ip>**). If setting a static IP, you must then identify the DNS server IP address (**nameserver=<dns_ip>**) on each node. The following example sets:

- The node's IP address to **10.10.10.2**
- The gateway address to **10.10.10.254**
- The netmask to **255.255.255.0**

- The hostname to **core0.example.com**
- The DNS server address to **4.4.4.41**
- The auto-configuration value to **none**. No auto-configuration is required when IP networking is configured statically.

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none
nameserver=4.4.4.41
```



NOTE

When you use DHCP to configure IP addressing for the RHCOS machines, the machines also obtain the DNS server information through DHCP. For DHCP-based deployments, you can define the DNS server address that is used by the RHCOS nodes through your DHCP server configuration.

Configuring an IP address without a static hostname

You can configure an IP address without assigning a static hostname. If a static hostname is not set by the user, it will be picked up and automatically set by a reverse DNS lookup. To configure an IP address without a static hostname refer to the following example:

- The node's IP address to **10.10.10.2**
- The gateway address to **10.10.10.254**
- The netmask to **255.255.255.0**
- The DNS server address to **4.4.4.41**
- The auto-configuration value to **none**. No auto-configuration is required when IP networking is configured statically.

```
ip=10.10.10.2::10.10.10.254:255.255.255.0::enp1s0:none
nameserver=4.4.4.41
```

Specifying multiple network interfaces

You can specify multiple network interfaces by setting multiple **ip=** entries.

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none
ip=10.10.10.3::10.10.10.254:255.255.255.0:core0.example.com:enp2s0:none
```

Configuring default gateway and route

Optional: You can configure routes to additional networks by setting an **rd.route=** value.



NOTE

When you configure one or multiple networks, one default gateway is required. If the additional network gateway is different from the primary network gateway, the default gateway must be the primary network gateway.

- Run the following command to configure the default gateway:


```
ip=::10.10.10.254:::
```

- Enter the following command to configure the route for the additional network:

```
rd.route=20.20.20.0/24:20.20.20.254:enp2s0
```

Disabling DHCP on a single interface

You can disable DHCP on a single interface, such as when there are two or more network interfaces and only one interface is being used. In the example, the **enp1s0** interface has a static networking configuration and DHCP is disabled for **enp2s0**, which is not used:

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none
ip=::::core0.example.com:enp2s0:none
```

Combining DHCP and static IP configurations

You can combine DHCP and static IP configurations on systems with multiple network interfaces, for example:

```
ip=enp1s0:dhcp
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp2s0:none
```

Configuring VLANs on individual interfaces

Optional: You can configure VLANs on individual interfaces by using the **vlan=** parameter.

- To configure a VLAN on a network interface and use a static IP address, run the following command:

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp2s0.100:none
vlan=enp2s0.100:enp2s0
```

- To configure a VLAN on a network interface and to use DHCP, run the following command:

```
ip=enp2s0.100:dhcp
vlan=enp2s0.100:enp2s0
```

Providing multiple DNS servers

You can provide multiple DNS servers by adding a **nameserver=** entry for each server, for example:

```
nameserver=1.1.1.1
nameserver=8.8.8.8
```

Bonding multiple network interfaces to a single interface

Optional: You can bond multiple network interfaces to a single interface by using the **bond=** option. Refer to the following examples:

- The syntax for configuring a bonded interface is: **bond=<name>[:<network_interfaces>][:options]**
<name> is the bonding device name (**bond0**), **<network_interfaces>** represents a comma-separated list of physical (ethernet) interfaces (**em1,em2**), and *options* is a comma-separated list of bonding options. Enter **modinfo bonding** to see available options.
- When you create a bonded interface using **bond=**, you must specify how the IP address is assigned and other information for the bonded interface.

- To configure the bonded interface to use DHCP, set the bond's IP address to **dhcp**. For example:

```
bond=bond0:em1,em2:mode=active-backup
ip=bond0:dhcp
```

- To configure the bonded interface to use a static IP address, enter the specific IP address you want and related information. For example:

```
bond=bond0:em1,em2:mode=active-backup
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:bond0:none
```

Bonding multiple SR-IOV network interfaces to a dual port NIC interface



IMPORTANT

Support for Day 1 operations associated with enabling NIC partitioning for SR-IOV devices is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

Optional: You can bond multiple SR-IOV network interfaces to a dual port NIC interface by using the **bond=** option.

On each node, you must perform the following tasks:

1. Create the SR-IOV virtual functions (VFs) following the guidance in [Managing SR-IOV devices](#). Follow the procedure in the "Attaching SR-IOV networking devices to virtual machines" section.
2. Create the bond, attach the desired VFs to the bond and set the bond link state up following the guidance in [Configuring network bonding](#). Follow any of the described procedures to create the bond.

The following examples illustrate the syntax you must use:

- The syntax for configuring a bonded interface is **bond=<name>[:<network_interfaces>][:options]**. **<name>** is the bonding device name (**bond0**), **<network_interfaces>** represents the virtual functions (VFs) by their known name in the kernel and shown in the output of the **ip link** command (**eno1f0**, **eno2f0**), and **options** is a comma-separated list of bonding options. Enter **modinfo bonding** to see available options.
- When you create a bonded interface using **bond=**, you must specify how the IP address is assigned and other information for the bonded interface.
 - To configure the bonded interface to use DHCP, set the bond's IP address to **dhcp**. For example:

```
bond=bond0:eno1f0,eno2f0:mode=active-backup
ip=bond0:dhcp
```

- o To configure the bonded interface to use a static IP address, enter the specific IP address you want and related information. For example:

```
bond=bond0:eno1f0,eno2f0:mode=active-backup
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:bond0:none
```

Using network teaming

Optional: You can use a network teaming as an alternative to bonding by using the **team=** parameter:

- The syntax for configuring a team interface is: **team=name[:network_interfaces]**
name is the team device name (**team0**) and *network_interfaces* represents a comma-separated list of physical (ethernet) interfaces (**em1, em2**).



NOTE

Teaming is planned to be deprecated when RHCOS switches to an upcoming version of RHEL. For more information, see this [Red Hat Knowledgebase Article](#).

Use the following example to configure a network team:

```
team=team0:em1,em2
ip=team0:dhcp
```


15.2.11.3.9.2. coreos-installer options for ISO and PXE installations

You can install RHCOS by running **coreos-installer install <options> <device>** at the command prompt, after booting into the RHCOS live environment from an ISO image.

The following table shows the subcommands, options, and arguments you can pass to the **coreos-installer** command.


Table 15.9. coreos-installer subcommands, command-line options, and arguments

coreos-installer install subcommand	
<i>Subcommand</i>	<i>Description</i>
\$ coreos-installer install <options> <device>	Embed an Ignition config in an ISO image.
coreos-installer install subcommand options	
<i>Option</i>	<i>Description</i>
-u, --image-url <url>	Specify the image URL manually.
-f, --image-file <path>	Specify a local image file manually. Used for debugging.
-i, --ignition-file <path>	Embed an Ignition config from a file.
-l, --ignition-url <URL>	Embed an Ignition config from a URL.

--ignition-hash <digest>	Digest type-value of the Ignition config.
-p, --platform <name>	Override the Ignition platform ID for the installed system.
--console <spec>	Set the kernel and bootloader console for the installed system. For more information about the format of <spec> , see the Linux kernel serial console documentation.
--append-karg <arg>...	Append a default kernel argument to the installed system.
--delete-karg <arg>...	Delete a default kernel argument from the installed system.
-n, --copy-network	Copy the network configuration from the install environment.  IMPORTANT The --copy-network option only copies networking configuration found under /etc/NetworkManager/system-connections . In particular, it does not copy the system hostname.
--network-dir <path>	For use with -n . Default is /etc/NetworkManager/system-connections/ .
--save-partlabel <lx>..	Save partitions with this label glob.
--save-partindex <id>...	Save partitions with this number or range.
--insecure	Skip RHCOS image signature verification.
--insecure-ignition	Allow Ignition URL without HTTPS or hash.
--architecture <name>	Target CPU architecture. Valid values are x86_64 and aarch64 .
--preserve-on-error	Do not clear partition table on error.
-h, --help	Print help information.
coreos-installer install subcommand argument	

<i>Argument</i>	<i>Description</i>
<device>	The destination device.
coreos-installer ISO subcommands	
<i>Subcommand</i>	<i>Description</i>
\$ coreos-installer iso customize <options> <ISO_image>	Customize a RHCOS live ISO image.
coreos-installer iso reset <options> <ISO_image>	Restore a RHCOS live ISO image to default settings.
coreos-installer iso ignition remove <options> <ISO_image>	Remove the embedded Ignition config from an ISO image.
coreos-installer ISO customize subcommand options	
<i>Option</i>	<i>Description</i>
--dest-ignition <path>	Merge the specified Ignition config file into a new configuration fragment for the destination system.
--dest-console <spec>	Specify the kernel and bootloader console for the destination system.
--dest-device <path>	Install and overwrite the specified destination device.
--dest-karg-append <arg>	Add a kernel argument to each boot of the destination system.
--dest-karg-delete <arg>	Delete a kernel argument from each boot of the destination system.
--network-keyfile <path>	Configure networking by using the specified NetworkManager keyfile for live and destination systems.
--ignition-ca <path>	Specify an additional TLS certificate authority to be trusted by Ignition.
--pre-install <path>	Run the specified script before installation.
--post-install <path>	Run the specified script after installation.
--installer-config <path>	Apply the specified installer configuration file.

--live-ignition <path>	Merge the specified Ignition config file into a new configuration fragment for the live environment.
--live-karg-append <arg>	Add a kernel argument to each boot of the live environment.
--live-karg-delete <arg>	Delete a kernel argument from each boot of the live environment.
--live-karg-replace <k=o=n>	Replace a kernel argument in each boot of the live environment, in the form key=old=new .
-f, --force	Overwrite an existing Ignition config.
-o, --output <path>	Write the ISO to a new output file.
-h, --help	Print help information.
coreos-installer PXE subcommands	
<i>Subcommand</i>	<i>Description</i>
Note that not all of these options are accepted by all subcommands.	
coreos-installer pxe customize <options> <path>	Customize a RHCOS live PXE boot config.
coreos-installer pxe ignition wrap <options>	Wrap an Ignition config in an image.
coreos-installer pxe ignition unwrap <options> <image_name>	Show the wrapped Ignition config in an image.
coreos-installer PXE customize subcommand options	
<i>Option</i>	<i>Description</i>
Note that not all of these options are accepted by all subcommands.	
--dest-ignition <path>	Merge the specified Ignition config file into a new configuration fragment for the destination system.
--dest-console <spec>	Specify the kernel and bootloader console for the destination system.
--dest-device <path>	Install and overwrite the specified destination device.

--network-keyfile <path>	Configure networking by using the specified NetworkManager keyfile for live and destination systems.
--ignition-ca <path>	Specify an additional TLS certificate authority to be trusted by Ignition.
--pre-install <path>	Run the specified script before installation.
post-install <path>	Run the specified script after installation.
--installer-config <path>	Apply the specified installer configuration file.
--live-ignition <path>	Merge the specified Ignition config file into a new configuration fragment for the live environment.
-o, --output <path>	Write the initramfs to a new output file.  NOTE This option is required for PXE environments.
-h, --help	Print help information.

15.2.11.3.9.3. coreos.inst boot options for ISO or PXE installations

You can automatically invoke **coreos-installer** options at boot time by passing **coreos.inst** boot arguments to the RHCOS live installer. These are provided in addition to the standard boot arguments.

- For ISO installations, the **coreos.inst** options can be added by interrupting the automatic boot at the bootloader menu. You can interrupt the automatic boot by pressing **TAB** while the **RHEL CoreOS (Live)** menu option is highlighted.
- For PXE or iPXE installations, the **coreos.inst** options must be added to the **APPEND** line before the RHCOS live installer is booted.

The following table shows the RHCOS live installer **coreos.inst** boot options for ISO and PXE installations.

Table 15.10. coreos.inst boot options

Argument	Description
coreos.inst.install_dev	Required. The block device on the system to install to. It is recommended to use the full path, such as /dev/sda , although sda is allowed.

Argument	Description
coreos.inst.ignition_url	Optional: The URL of the Ignition config to embed into the installed system. If no URL is specified, no Ignition config is embedded. Only HTTP and HTTPS protocols are supported.
coreos.inst.save_partlabel	Optional: Comma-separated labels of partitions to preserve during the install. Glob-style wildcards are permitted. The specified partitions do not need to exist.
coreos.inst.save_partindex	Optional: Comma-separated indexes of partitions to preserve during the install. Ranges m-n are permitted, and either m or n can be omitted. The specified partitions do not need to exist.
coreos.inst.insecure	Optional: Permits the OS image that is specified by coreos.inst.image_url to be unsigned.
coreos.inst.image_url	<p>Optional: Download and install the specified RHCOS image.</p> <ul style="list-style-type: none"> ● This argument should not be used in production environments and is intended for debugging purposes only. ● While this argument can be used to install a version of RHCOS that does not match the live media, it is recommended that you instead use the media that matches the version you want to install. ● If you are using coreos.inst.image_url, you must also use coreos.inst.insecure. This is because the bare-metal media are not GPG-signed for OpenShift Container Platform. ● Only HTTP and HTTPS protocols are supported.
coreos.inst.skip_reboot	Optional: The system will not reboot after installing. After the install finishes, you will receive a prompt that allows you to inspect what is happening during installation. This argument should not be used in production environments and is intended for debugging purposes only.


Argument	Description
coreos.inst.platform_id	Optional: The Ignition platform ID of the platform the RHCOS image is being installed on. Default is metal . This option determines whether or not to request an Ignition config from the cloud provider, such as VMware. For example: coreos.inst.platform_id=vmware.
ignition.config.url	Optional: The URL of the Ignition config for the live boot. For example, this can be used to customize how coreos-installer is invoked, or to run code before or after the installation. This is different from coreos.inst.ignition_url , which is the Ignition config for the installed system.

15.2.11.4. Enabling multipathing with kernel arguments on RHCOS

RHCOS supports multipathing on the primary disk, allowing stronger resilience to hardware failure to achieve higher host availability.

You can enable multipathing at installation time for nodes that were provisioned in OpenShift Container Platform 4.8 or later. While postinstallation support is available by activating multipathing via the machine config, enabling multipathing during installation is recommended.


In setups where any I/O to non-optimized paths results in I/O system errors, you must enable multipathing at installation time.



IMPORTANT

On IBM Z® and IBM® LinuxONE, you can enable multipathing only if you configured your cluster for it during installation. For more information, see "Installing RHCOS and starting the OpenShift Container Platform bootstrap process" in *Installing a cluster with z/VM on IBM Z® and IBM® LinuxONE*.

The following procedure enables multipath at installation time and appends kernel arguments to the **coreos-installer install** command so that the installed system itself will use multipath beginning from the first boot.



NOTE

OpenShift Container Platform does not support enabling multipathing as a day-2 activity on nodes that have been upgraded from 4.6 or earlier.

Procedure

1. To enable multipath and start the **multipathd** daemon, run the following command on the installation host:

```
$ mpathconf --enable && systemctl start multipathd.service
```

- Optional: If booting the PXE or ISO, you can instead enable multipath by adding **rd.multipath=default** from the kernel command line.
2. Append the kernel arguments by invoking the **coreos-installer** program:
 - If there is only one multipath device connected to the machine, it should be available at path **/dev/mapper/mpatha**. For example:

```
$ coreos-installer install /dev/mapper/mpatha \ 1
--append-karg rd.multipath=default \
--append-karg root=/dev/disk/by-label/dm-mpath-root \
--append-karg rw
```

1 Indicates the path of the single multipathed device.

- If there are multiple multipath devices connected to the machine, or to be more explicit, instead of using **/dev/mapper/mpatha**, it is recommended to use the World Wide Name (WWN) symlink available in **/dev/disk/by-id**. For example:

```
$ coreos-installer install /dev/disk/by-id/wwn-<wwn_ID> \ 1
--append-karg rd.multipath=default \
--append-karg root=/dev/disk/by-label/dm-mpath-root \
--append-karg rw
```

1 Indicates the WWN ID of the target multipathed device. For example, **0xx194e957fcedb4841**.

This symlink can also be used as the **coreos.inst.install_dev** kernel argument when using special **coreos.inst.*** arguments to direct the live installer. For more information, see "Installing RHCOS and starting the OpenShift Container Platform bootstrap process".

3. Check that the kernel arguments worked by going to one of the worker nodes and listing the kernel command line arguments (in **/proc/cmdline** on the host):

```
$ oc debug node/ip-10-0-141-105.ec2.internal
```

Example output

```
Starting pod/ip-10-0-141-105ec2internal-debug ...
To use host binaries, run `chroot /host`

sh-4.2# cat /host/proc/cmdline
...
rd.multipath=default root=/dev/disk/by-label/dm-mpath-root
...

sh-4.2# exit
```

You should see the added kernel arguments.

15.2.11.5. Installing RHCOS manually on an iSCSI boot device

You can manually install RHCOS on an iSCSI target.

Prerequisites

1. You are in the RHCOS live environment.
2. You have an iSCSI target that you want to install RHCOS on.

Procedure

1. Mount the iSCSI target from the live environment by running the following command:

```
$ iscsiadm \
  --mode discovery \
  --type sendtargets \
  --portal <IP_address> \ 1
  --login
```

- 1 The IP address of the target portal.

2. Install RHCOS onto the iSCSI target by running the following command and using the necessary kernel arguments, for example:

```
$ coreos-installer install \
  /dev/disk/by-path/ip-<IP_address>:<port>-iscsi-<target_iqn>-lun-<lun> \ 1
  --append-karg rd.iscsi.initiator=<initiator_iqn> \ 2
  --append.karg netroot=<target_iqn> \ 3
  --console ttyS0,115200n8
  --ignition-file <path_to_file>
```

- 1 The location you are installing to. You must provide the IP address of the target portal, the associated port number, the target iSCSI node in IQN format, and the iSCSI logical unit number (LUN).
- 2 The iSCSI initiator, or client, name in IQN format. The initiator forms a session to connect to the iSCSI target.
- 3 The the iSCSI target, or server, name in IQN format.

For more information about the iSCSI options supported by **dracut**, see the [dracut.cmdline manual page](#).

3. Unmount the iSCSI disk with the following command:

```
$ iscsiadm --mode node --logoutall=all
```

This procedure can also be performed using the **coreos-installer iso customize** or **coreos-installer pxe customize** subcommands.

15.2.11.6. Installing RHCOS on an iSCSI boot device using iBFT

On a completely diskless machine, the iSCSI target and initiator values can be passed through iBFT. iSCSI multipathing is also supported.

Prerequisites

1. You are in the RHCOS live environment.
2. You have an iSCSI target you want to install RHCOS on.
3. Optional: you have multipathed your iSCSI target.

Procedure

1. Mount the iSCSI target from the live environment by running the following command:

```
$ iscsiadm \  
  --mode discovery \  
  --type sendtargets \  
  --portal <IP_address> \ 1 \  
  --login
```

- 1** The IP address of the target portal.

2. Optional: enable multipathing and start the daemon with the following command:

```
$ mpathconf --enable && systemctl start multipathd.service
```

3. Install RHCOS onto the iSCSI target by running the following command and using the necessary kernel arguments, for example:

```
$ coreos-installer install \  
  /dev/mapper/mpatha \ 1 \  
  --append-karg rd.iscsi.firmware=1 \ 2 \  
  --append-karg rd.multipath=default \ 3 \  
  --console ttyS0 \  
  --ignition-file <path_to_file>
```

- 1** The path of a single multipathed device. If there are multiple multipath devices connected, or to be explicit, you can use the World Wide Name (WWN) symlink available in **/dev/disk/by-path**.
- 2** The iSCSI parameter is read from the BIOS firmware.
- 3** Optional: include this parameter if you are enabling multipathing.

For more information about the iSCSI options supported by **dracut**, see the [dracut.cmdline manual page](#).

4. Unmount the iSCSI disk:

```
$ iscsiadm --mode node --logout=all
```

This procedure can also be performed using the **coreos-installer iso customize** or **coreos-installer pxe customize** subcommands.

Additional resources

- See [Installing RHCOS and starting the OpenShift Container Platform bootstrap process](#) for more information on using special **coreos.inst.*** arguments to direct the live installer.

15.2.12. Waiting for the bootstrap process to complete

The OpenShift Container Platform bootstrap process begins after the cluster nodes first boot into the persistent RHCOS environment that has been installed to disk. The configuration information provided through the Ignition config files is used to initialize the bootstrap process and install OpenShift Container Platform on the machines. You must wait for the bootstrap process to complete.

Prerequisites

- You have created the Ignition config files for your cluster.
- You have configured suitable network, DNS and load balancing infrastructure.
- You have obtained the installation program and generated the Ignition config files for your cluster.
- You installed RHCOS on your cluster machines and provided the Ignition config files that the OpenShift Container Platform installation program generated.
- Your machines have direct internet access or have an HTTP or HTTPS proxy available.

Procedure

1. Monitor the bootstrap process:

```
$ ./openshift-install --dir <installation_directory> wait-for bootstrap-complete \ 1
--log-level=info 2
```

- 1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.


- 2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Example output

```
INFO Waiting up to 30m0s for the Kubernetes API at https://api.test.example.com:6443...
INFO API v1.29.4 up
INFO Waiting up to 30m0s for bootstrapping to complete...
INFO It is now safe to remove the bootstrap resources
```

The command succeeds when the Kubernetes API server signals that it has been bootstrapped on the control plane machines.

2. After the bootstrap process is complete, remove the bootstrap machine from the load balancer.



IMPORTANT

You must remove the bootstrap machine from the load balancer at this point. You can also remove or reformat the bootstrap machine itself.

Additional resources

- See [Monitoring installation progress](#) for more information about monitoring the installation logs and retrieving diagnostic data if installation issues arise.

15.2.13. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

15.2.14. Approving the certificate signing requests for your machines

When you add machines to a cluster, two pending certificate signing requests (CSRs) are generated for each machine that you added. You must confirm that these CSRs are approved or, if necessary, approve them yourself. The client requests must be approved first, followed by the server requests.

Prerequisites

- You added machines to your cluster.

Procedure

1. Confirm that the cluster recognizes the machines:

```
$ oc get nodes
```

Example output

NAME	STATUS	ROLES	AGE	VERSION
master-0	Ready	master	63m	v1.29.4
master-1	Ready	master	63m	v1.29.4
master-2	Ready	master	64m	v1.29.4

The output lists all of the machines that you created.



NOTE

The preceding output might not include the compute nodes, also known as worker nodes, until some CSRs are approved.

- Review the pending CSRs and ensure that you see the client requests with the **Pending** or **Approved** status for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

NAME	AGE	REQUESTOR	CONDITION
csr-8b2br	15m	system:serviceaccount:openshift-machine-config-operator:node-bootstrapper	Pending
csr-8vnps	15m	system:serviceaccount:openshift-machine-config-operator:node-bootstrapper	Pending
...			

In this example, two machines are joining the cluster. You might see more approved CSRs in the list.

- If the CSRs were not approved, after all of the pending CSRs for the machines you added are in **Pending** status, approve the CSRs for your cluster machines:



NOTE

Because the CSRs rotate automatically, approve your CSRs within an hour of adding the machines to the cluster. If you do not approve them within an hour, the certificates will rotate, and more than two certificates will be present for each node. You must approve all of these certificates. After the client CSR is approved, the Kubelet creates a secondary CSR for the serving certificate, which requires manual approval. Then, subsequent serving certificate renewal requests are automatically approved by the **machine-approver** if the Kubelet requests a new certificate with identical parameters.

**NOTE**

For clusters running on platforms that are not machine API enabled, such as bare metal and other user-provisioned infrastructure, you must implement a method of automatically approving the kubelet serving certificate requests (CSRs). If a request is not approved, then the **oc exec**, **oc rsh**, and **oc logs** commands cannot succeed, because a serving certificate is required when the API server connects to the kubelet. Any operation that contacts the Kubelet endpoint requires this certificate approval to be in place. The method must watch for new CSRs, confirm that the CSR was submitted by the **node-bootstrap** service account in the **system:node** or **system:admin** groups, and confirm the identity of the node.

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{"\n"}{{end}}{{end}}' | xargs --no-run-if-empty oc adm certificate approve
```

**NOTE**

Some Operators might not become available until some CSRs are approved.

- Now that your client requests are approved, you must review the server requests for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

```
NAME      AGE   REQUESTOR                                     CONDITION
csr-bfd72 5m26s system:node:ip-10-0-50-126.us-east-2.compute.internal
Pending
csr-c57lv 5m26s system:node:ip-10-0-95-157.us-east-2.compute.internal
Pending
...
```

- If the remaining CSRs are not approved, and are in the **Pending** status, approve the CSRs for your cluster machines:

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{\n"}\n{{end}}\n' | xargs oc adm certificate approve
```

6. After all client and server CSRs have been approved, the machines have the **Ready** status. Verify this by running the following command:

```
$ oc get nodes
```

Example output

```
NAME      STATUS  ROLES  AGE  VERSION
master-0  Ready   master 73m  v1.29.4
master-1  Ready   master 73m  v1.29.4
master-2  Ready   master 74m  v1.29.4
worker-0  Ready   worker 11m  v1.29.4
worker-1  Ready   worker 11m  v1.29.4
```



NOTE

It can take a few minutes after approval of the server CSRs for the machines to transition to the **Ready** status.

Additional information

- For more information on CSRs, see [Certificate Signing Requests](#).

15.2.15. Initial Operator configuration

After the control plane initializes, you must immediately configure some Operators so that they all become available.

Prerequisites

- Your control plane has initialized.

Procedure

1. Watch the cluster components come online:

```
$ watch -n5 oc get clusteroperators
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.16.0	True	False	False	19m
baremetal	4.16.0	True	False	False	37m
cloud-credential	4.16.0	True	False	False	40m
cluster-autoscaler	4.16.0	True	False	False	37m
config-operator	4.16.0	True	False	False	38m

console	4.16.0	True	False	False	26m
csi-snapshot-controller	4.16.0	True	False	False	37m
dns	4.16.0	True	False	False	37m
etcd	4.16.0	True	False	False	36m
image-registry	4.16.0	True	False	False	31m
ingress	4.16.0	True	False	False	30m
insights	4.16.0	True	False	False	31m
kube-apiserver	4.16.0	True	False	False	26m
kube-controller-manager	4.16.0	True	False	False	36m
kube-scheduler	4.16.0	True	False	False	36m
kube-storage-version-migrator	4.16.0	True	False	False	37m
machine-api	4.16.0	True	False	False	29m
machine-approver	4.16.0	True	False	False	37m
machine-config	4.16.0	True	False	False	36m
marketplace	4.16.0	True	False	False	37m
monitoring	4.16.0	True	False	False	29m
network	4.16.0	True	False	False	38m
node-tuning	4.16.0	True	False	False	37m
openshift-apiserver	4.16.0	True	False	False	32m
openshift-controller-manager	4.16.0	True	False	False	30m
openshift-samples	4.16.0	True	False	False	32m
operator-lifecycle-manager	4.16.0	True	False	False	37m
operator-lifecycle-manager-catalog	4.16.0	True	False	False	37m
operator-lifecycle-manager-packageserver	4.16.0	True	False	False	32m
service-ca	4.16.0	True	False	False	38m
storage	4.16.0	True	False	False	37m

2. Configure the Operators that are not available.

Additional resources

- See [Gathering logs from a failed installation](#) for details about gathering data in the event of a failed OpenShift Container Platform installation.
- See [Troubleshooting Operator issues](#) for steps to check Operator pod health across the cluster and gather Operator logs for diagnosis.

15.2.15.1. Image registry removed during installation

On platforms that do not provide shareable object storage, the OpenShift Image Registry Operator bootstraps itself as **Removed**. This allows **openshift-installer** to complete installations on these platform types.

After installation, you must edit the Image Registry Operator configuration to switch the **managementState** from **Removed** to **Managed**. When this has completed, you must configure storage.

15.2.15.2. Image registry storage configuration

The Image Registry Operator is not initially available for platforms that do not provide default storage. After installation, you must configure your registry to use storage so that the Registry Operator is made available.

Instructions are shown for configuring a persistent volume, which is required for production clusters. Where applicable, instructions are shown for configuring an empty directory as the storage location, which is available for only non-production clusters.


Additional instructions are provided for allowing the image registry to use block storage types by using the **Recreate** rollout strategy during upgrades.

15.2.15.2.1. Configuring registry storage for bare metal and other manual installations

As a cluster administrator, following installation you must configure your registry to use storage.

Prerequisites

- You have access to the cluster as a user with the **cluster-admin** role.
- You have a cluster that uses manually-provisioned Red Hat Enterprise Linux CoreOS (RHCOS) nodes, such as bare metal.
- You have provisioned persistent storage for your cluster, such as Red Hat OpenShift Data Foundation.



IMPORTANT

OpenShift Container Platform supports **ReadWriteOnce** access for image registry storage when you have only one replica. **ReadWriteOnce** access also requires that the registry uses the **Recreate** rollout strategy. To deploy an image registry that supports high availability with two or more replicas, **ReadWriteMany** access is required.

- Must have 100Gi capacity.

Procedure

1. To configure your registry to use storage, change the **spec.storage.pvc** in the **configs.imageregistry/cluster** resource.



NOTE

When you use shared storage, review your security settings to prevent outside access.

2. Verify that you do not have a registry pod:

```
$ oc get pod -n openshift-image-registry -l docker-registry=default
```

Example output

```
No resources found in openshift-image-registry namespace
```



NOTE

If you do have a registry pod in your output, you do not need to continue with this procedure.

3. Check the registry configuration:

```
$ oc edit configs.imageregistry.operator.openshift.io
```

Example output

```
storage:
  pvc:
    claim:
```

Leave the **claim** field blank to allow the automatic creation of an **image-registry-storage** PVC.

4. Check the **clusteroperator** status:

```
$ oc get clusteroperator image-registry
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
image-registry	4.16	True	False	False	6h50m

5. Ensure that your registry is set to managed to enable building and pushing of images.

- Run:

```
$ oc edit configs.imageregistry/cluster
```

Then, change the line

```
managementState: Removed
```

to

```
managementState: Managed
```

15.2.15.2.2. Configuring storage for the image registry in non-production clusters

You must configure storage for the Image Registry Operator. For non-production clusters, you can set the image registry to an empty directory. If you do so, all images are lost if you restart the registry.

Procedure

- To set the image registry storage to an empty directory:

```
$ oc patch configs.imageregistry.operator.openshift.io cluster --type merge --patch '{"spec": {"storage":{"emptyDir":{}}}'
```

**WARNING**

Configure this option for only non-production clusters.


If you run this command before the Image Registry Operator initializes its components, the **oc patch** command fails with the following error:

```
Error from server (NotFound): configs.imageregistry.operator.openshift.io "cluster" not found
```

Wait a few minutes and run the command again.

15.2.15.2.3. Configuring block registry storage for bare metal

To allow the image registry to use block storage types during upgrades as a cluster administrator, you can use the **Recreate** rollout strategy.

**IMPORTANT**

Block storage volumes, or block persistent volumes, are supported but not recommended for use with the image registry on production clusters. An installation where the registry is configured on block storage is not highly available because the registry cannot have more than one replica.

If you choose to use a block storage volume with the image registry, you must use a filesystem persistent volume claim (PVC).

Procedure

1. Enter the following command to set the image registry storage as a block storage type, patch the registry so that it uses the **Recreate** rollout strategy, and runs with only one (**1**) replica:

```
$ oc patch config.imageregistry.operator.openshift.io/cluster --type=merge -p '{"spec": {"rolloutStrategy": "Recreate", "replicas": 1}}'
```

2. Provision the PV for the block storage device, and create a PVC for that volume. The requested block volume uses the ReadWriteOnce (RWO) access mode.
 - a. Create a **pvc.yaml** file with the following contents to define a VMware vSphere **PersistentVolumeClaim** object:

```
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: image-registry-storage 1
  namespace: openshift-image-registry 2
spec:
  accessModes:
    - ReadWriteOnce 3
```

```
resources:
  requests:
    storage: 100Gi 4
```

- 1** A unique name that represents the **PersistentVolumeClaim** object.
- 2** The namespace for the **PersistentVolumeClaim** object, which is **openshift-image-registry**.
- 3** The access mode of the persistent volume claim. With **ReadWriteOnce**, the volume can be mounted with read and write permissions by a single node.
- 4** The size of the persistent volume claim.

b. Enter the following command to create the **PersistentVolumeClaim** object from the file:

```
$ oc create -f pvc.yaml -n openshift-image-registry
```

3. Enter the following command to edit the registry configuration so that it references the correct PVC:

```
$ oc edit config.imageregistry.operator.openshift.io -o yaml
```

Example output

```
storage:
  pvc:
    claim: 1
```

- 1** By creating a custom PVC, you can leave the **claim** field blank for the default automatic creation of an **image-registry-storage** PVC.

15.2.16. Completing installation on user-provisioned infrastructure

After you complete the Operator configuration, you can finish installing the cluster on infrastructure that you provide.

Prerequisites

- Your control plane has initialized.
- You have completed the initial Operator configuration.

Procedure

1. Confirm that all the cluster components are online with the following command:

```
$ watch -n5 oc get clusteroperators
```

Example output

NAME SINCE	VERSION	AVAILABLE	PROGRESSING	DEGRADED	
authentication	4.16.0	True	False	False	19m
baremetal	4.16.0	True	False	False	37m
cloud-credential	4.16.0	True	False	False	40m
cluster-autoscaler	4.16.0	True	False	False	37m
config-operator	4.16.0	True	False	False	38m
console	4.16.0	True	False	False	26m
csi-snapshot-controller	4.16.0	True	False	False	37m
dns	4.16.0	True	False	False	37m
etcd	4.16.0	True	False	False	36m
image-registry	4.16.0	True	False	False	31m
ingress	4.16.0	True	False	False	30m
insights	4.16.0	True	False	False	31m
kube-apiserver	4.16.0	True	False	False	26m
kube-controller-manager	4.16.0	True	False	False	36m
kube-scheduler	4.16.0	True	False	False	36m
kube-storage-version-migrator	4.16.0	True	False	False	37m
machine-api	4.16.0	True	False	False	29m
machine-approver	4.16.0	True	False	False	37m
machine-config	4.16.0	True	False	False	36m
marketplace	4.16.0	True	False	False	37m
monitoring	4.16.0	True	False	False	29m
network	4.16.0	True	False	False	38m
node-tuning	4.16.0	True	False	False	37m
openshift-apiserver	4.16.0	True	False	False	32m
openshift-controller-manager	4.16.0	True	False	False	30m
openshift-samples	4.16.0	True	False	False	32m
operator-lifecycle-manager	4.16.0	True	False	False	37m
operator-lifecycle-manager-catalog	4.16.0	True	False	False	37m
operator-lifecycle-manager-packageserver	4.16.0	True	False	False	32m
service-ca	4.16.0	True	False	False	38m
storage	4.16.0	True	False	False	37m

Alternatively, the following command notifies you when all of the clusters are available. It also retrieves and displays credentials:

```
$ ./openshift-install --dir <installation_directory> wait-for install-complete 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

Example output

```
INFO Waiting up to 30m0s for the cluster to initialize...
```

The command succeeds when the Cluster Version Operator finishes deploying the OpenShift Container Platform cluster from Kubernetes API server.



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

2. Confirm that the Kubernetes API server is communicating with the pods.

- a. To view a list of all pods, use the following command:

```
$ oc get pods --all-namespaces
```

Example output

```

NAMESPACE           NAME                                     READY  STATUS
RESTARTS  AGE
openshift-apiserver-operator  openshift-apiserver-operator-85cb746d55-zqhs8  1/1
Running   1      9m
openshift-apiserver          apiserver-67b9g                                1/1  Running  0
3m
openshift-apiserver          apiserver-ljcmx                                1/1  Running  0
1m
openshift-apiserver          apiserver-z25h4                                1/1  Running  0
2m
openshift-authentication-operator  authentication-operator-69d5d8bf84-vh2n8      1/1
Running   0      5m
...

```

- b. View the logs for a pod that is listed in the output of the previous command by using the following command:

```
$ oc logs <pod_name> -n <namespace> ❶
```

- ❶ Specify the pod name and namespace, as shown in the output of the previous command.

If the pod logs display, the Kubernetes API server can communicate with the cluster machines.

3. For an installation with Fibre Channel Protocol (FCP), additional steps are required to enable multipathing. Do not enable multipathing during installation. See "Enabling multipathing with kernel arguments on RHCOS" in the *Postinstallation machine configuration tasks* documentation for more information.

15.2.17. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service

15.2.18. Next steps

- [Validating an installation.](#)
- [Customize your cluster.](#)
- If necessary, you can [opt out of remote health reporting](#) .
- [Set up your registry and configure registry storage](#) .

15.3. INSTALLING A USER-PROVISIONED BARE METAL CLUSTER WITH NETWORK CUSTOMIZATIONS

In OpenShift Container Platform 4.16, you can install a cluster on bare metal infrastructure that you provision with customized network configuration options. By customizing your network configuration, your cluster can coexist with existing IP address allocations in your environment and integrate with existing MTU and VXLAN configurations.

When you customize OpenShift Container Platform networking, you must set most of the network configuration parameters during installation. You can modify only **kubeProxy** network configuration parameters in a running cluster.

15.3.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- If you use a firewall and plan to use the Telemetry service, you [configured the firewall to allow the sites](#) that your cluster requires access to.

15.3.2. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.



IMPORTANT

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

Additional resources

- See [Installing a user-provisioned bare metal cluster on a restricted network](#) for more information about performing a restricted network installation on bare metal infrastructure that you provision.

15.3.3. Requirements for a cluster with user-provisioned infrastructure

For a cluster that contains user-provisioned infrastructure, you must deploy all of the required machines.

This section describes the requirements for deploying OpenShift Container Platform on user-provisioned infrastructure.

15.3.3.1. Required machines for cluster installation

The smallest OpenShift Container Platform clusters require the following hosts:

Table 15.11. Minimum required hosts

Hosts	Description
One temporary bootstrap machine	The cluster requires the bootstrap machine to deploy the OpenShift Container Platform cluster on the three control plane machines. You can remove the bootstrap machine after you install the cluster.
Three control plane machines	The control plane machines run the Kubernetes and OpenShift Container Platform services that form the control plane.
At least two compute machines, which are also known as worker machines.	The workloads requested by OpenShift Container Platform users run on the compute machines.

**NOTE**

As an exception, you can run zero compute machines in a bare metal cluster that consists of three control plane machines only. This provides smaller, more resource efficient clusters for cluster administrators and developers to use for testing, development, and production. Running one compute machine is not supported.

**IMPORTANT**

To maintain high availability of your cluster, use separate physical hosts for these cluster machines.

The bootstrap and control plane machines must use Red Hat Enterprise Linux CoreOS (RHCOS) as the operating system. However, the compute machines can choose between Red Hat Enterprise Linux CoreOS (RHCOS), Red Hat Enterprise Linux (RHEL) 8.6 and later.

Note that RHCOS is based on Red Hat Enterprise Linux (RHEL) 9.2 and inherits all of its hardware certifications and requirements. See [Red Hat Enterprise Linux technology capabilities and limits](#) .

15.3.3.2. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 15.12. Minimum resource requirements

Machine	Operating System	CPU [1]	RAM	Storage	Input/Output Per Second (IOPS)[2]
Bootstrap	RHCOS	4	16 GB	100 GB	300
Control plane	RHCOS	4	16 GB	100 GB	300
Compute	RHCOS, RHEL 8.6 and later [3]	2	8 GB	100 GB	300

1. One CPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio: (threads per core × cores) × sockets = CPUs.
2. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes which require a 10 ms p99 fsync duration. Note that on many cloud platforms, storage size and IOPS scale together, so you might need to over-allocate storage volume to obtain sufficient performance.
3. As with all user-provisioned installations, if you choose to use RHEL compute machines in your cluster, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Use of RHEL 7 compute machines is deprecated and has been removed in OpenShift Container Platform 4.10 and later.



NOTE

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

15.3.3.3. Certificate signing requests management

Because your cluster has limited access to automatic machine management when you use infrastructure that you provision, you must provide a mechanism for approving cluster certificate signing requests (CSRs) after installation. The **kube-controller-manager** only approves the kubelet client CSRs. The **machine-approver** cannot guarantee the validity of a serving certificate that is requested by using kubelet credentials because it cannot confirm that the correct machine issued the request. You must determine and implement a method of verifying the validity of the kubelet serving certificate requests and approving them.

Additional resources

- See [Configuring a three-node cluster](#) for details about deploying three-node clusters in bare metal environments.
- See [Approving the certificate signing requests for your machines](#) for more information about approving cluster certificate signing requests after installation.

15.3.3.4. Networking requirements for user-provisioned infrastructure

All the Red Hat Enterprise Linux CoreOS (RHCOS) machines require networking to be configured in **initramfs** during boot to fetch their Ignition config files.

During the initial boot, the machines require an IP address configuration that is set either through a DHCP server or statically by providing the required boot options. After a network connection is established, the machines download their Ignition config files from an HTTP or HTTPS server. The Ignition config files are then used to set the exact state of each machine. The Machine Config Operator completes more changes to the machines, such as the application of new certificates or keys, after installation.

It is recommended to use a DHCP server for long-term management of the cluster machines. Ensure that the DHCP server is configured to provide persistent IP addresses, DNS server information, and hostnames to the cluster machines.



NOTE

If a DHCP service is not available for your user-provisioned infrastructure, you can instead provide the IP networking configuration and the address of the DNS server to the nodes at RHCOS install time. These can be passed as boot arguments if you are installing from an ISO image. See the *Installing RHCOS and starting the OpenShift Container Platform bootstrap process* section for more information about static IP provisioning and advanced networking options.

The Kubernetes API server must be able to resolve the node names of the cluster machines. If the API servers and worker nodes are in different zones, you can configure a default DNS search zone to allow the API server to resolve the node names. Another supported approach is to always refer to hosts by their fully-qualified domain names in both the node objects and all DNS requests.

15.3.3.4.1. Setting the cluster node hostnames through DHCP


On Red Hat Enterprise Linux CoreOS (RHCOS) machines, the hostname is set through NetworkManager. By default, the machines obtain their hostname through DHCP. If the hostname is not provided by DHCP, set statically through kernel arguments, or another method, it is obtained through a reverse DNS lookup. Reverse DNS lookup occurs after the network has been initialized on a node and can take time to resolve. Other system services can start prior to this and detect the hostname as **localhost** or similar. You can avoid this by using DHCP to provide the hostname for each cluster node.

Additionally, setting the hostnames through DHCP can bypass any manual DNS record name configuration errors in environments that have a DNS split-horizon implementation.

15.3.3.4.2. Network connectivity requirements

You must configure the network connectivity between machines to allow OpenShift Container Platform cluster components to communicate. Each machine must be able to resolve the hostnames of all other machines in the cluster.

This section provides details about the ports that are required.



IMPORTANT

In connected OpenShift Container Platform environments, all nodes are required to have internet access to pull images for platform containers and provide telemetry data to Red Hat.

Table 15.13. Ports used for all-machine to all-machine communications

Protocol	Port	Description
ICMP	N/A	Network reachability tests
TCP	1936	Metrics

Protocol	Port	Description
	9000-9999	Host level services, including the node exporter on ports 9100-9101 and the Cluster Version Operator on port 9099 .
	10250-10259	The default ports that Kubernetes reserves
UDP	4789	VXLAN
	6081	Geneve
	9000-9999	Host level services, including the node exporter on ports 9100-9101 .
	500	IPsec IKE packets
	4500	IPsec NAT-T packets
	123	Network Time Protocol (NTP) on UDP port 123 If an external NTP time server is configured, you must open UDP port 123 .
TCP/UDP	30000-32767	Kubernetes node port
ESP	N/A	IPsec Encapsulating Security Payload (ESP)

Table 15.14. Ports used for all-machine to control plane communications

Protocol	Port	Description
TCP	6443	Kubernetes API

Table 15.15. Ports used for control plane machine to control plane machine communications

Protocol	Port	Description
TCP	2379-2380	etcd server and peer ports

NTP configuration for user-provisioned infrastructure

OpenShift Container Platform clusters are configured to use a public Network Time Protocol (NTP) server by default. If you want to use a local enterprise NTP server, or if your cluster is being deployed in a disconnected network, you can configure the cluster to use a specific time server. For more information, see the documentation for *Configuring chrony time service*.

If a DHCP server provides NTP server information, the chrony time service on the Red Hat Enterprise Linux CoreOS (RHCOS) machines read the information and can sync the clock with the NTP servers.

Additional resources

- [Configuring chrony time service](#)


15.3.3.5. User-provisioned DNS requirements

In OpenShift Container Platform deployments, DNS name resolution is required for the following components:

- The Kubernetes API
- The OpenShift Container Platform application wildcard
- The bootstrap, control plane, and compute machines

Reverse DNS resolution is also required for the Kubernetes API, the bootstrap machine, the control plane machines, and the compute machines.

DNS A/AAAA or CNAME records are used for name resolution and PTR records are used for reverse name resolution. The reverse records are important because Red Hat Enterprise Linux CoreOS (RHCOS) uses the reverse records to set the hostnames for all the nodes, unless the hostnames are provided by DHCP. Additionally, the reverse records are used to generate the certificate signing requests (CSR) that OpenShift Container Platform needs to operate.




NOTE


It is recommended to use a DHCP server to provide the hostnames to each cluster node. See the *DHCP recommendations for user-provisioned infrastructure* section for more information.

The following DNS records are required for a user-provisioned OpenShift Container Platform cluster and they must be in place before installation. In each record, **<cluster_name>** is the cluster name and **<base_domain>** is the base domain that you specify in the **install-config.yaml** file. A complete DNS record takes the form: **<component>.<cluster_name>.<base_domain>..**

Table 15.16. Required DNS records

Component	Record	Description
Kubernetes API	api.<cluster_name>.<base_domain>..	A DNS A/AAAA or CNAME record, and a DNS PTR record, to identify the API load balancer. These records must be resolvable by both clients external to the cluster and from all the nodes within the cluster.

Component	Record	Description
	api-int.<cluster_name>.<base_domain>.	<p>A DNS A/AAAA or CNAME record, and a DNS PTR record, to internally identify the API load balancer. These records must be resolvable from all the nodes within the cluster.</p> <div style="display: flex; align-items: flex-start;">  <div> <p>IMPORTANT</p> <p>The API server must be able to resolve the worker nodes by the hostnames that are recorded in Kubernetes. If the API server cannot resolve the node names, then proxied API calls can fail, and you cannot retrieve logs from pods.</p> </div> </div>
Routes	*.apps.<cluster_name>.<base_domain>.	<p>A wildcard DNS A/AAAA or CNAME record that refers to the application ingress load balancer. The application ingress load balancer targets the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default. These records must be resolvable by both clients external to the cluster and from all the nodes within the cluster.</p> <p>For example, console-openshift-console.apps.<cluster_name>.<base_domain> is used as a wildcard route to the OpenShift Container Platform console.</p>
Bootstrap machine	bootstrap.<cluster_name>.<base_domain>.	A DNS A/AAAA or CNAME record, and a DNS PTR record, to identify the bootstrap machine. These records must be resolvable by the nodes within the cluster.
Control plane machines	<control_plane><n>.<cluster_name>.<base_domain>.	DNS A/AAAA or CNAME records and DNS PTR records to identify each machine for the control plane nodes. These records must be resolvable by the nodes within the cluster.
Compute machines	<compute><n>.<cluster_name>.<base_domain>.	DNS A/AAAA or CNAME records and DNS PTR records to identify each machine for the worker nodes. These records must be resolvable by the nodes within the cluster.



NOTE

In OpenShift Container Platform 4.4 and later, you do not need to specify etcd host and SRV records in your DNS configuration.

TIP

You can use the **dig** command to verify name and reverse name resolution. See the section on *Validating DNS resolution for user-provisioned infrastructure* for detailed validation steps.

15.3.3.5.1. Example DNS configuration for user-provisioned clusters

This section provides A and PTR record configuration samples that meet the DNS requirements for deploying OpenShift Container Platform on user-provisioned infrastructure. The samples are not meant to provide advice for choosing one DNS solution over another.

In the examples, the cluster name is **ocp4** and the base domain is **example.com**.

Example DNS A record configuration for a user-provisioned cluster

The following example is a BIND zone file that shows sample A records for name resolution in a user-provisioned cluster.

Example 15.4. Sample DNS zone database

```
$TTL 1W
@ IN SOA ns1.example.com. root (
    2019070700 ; serial
    3H ; refresh (3 hours)
    30M ; retry (30 minutes)
    2W ; expiry (2 weeks)
    1W ) ; minimum (1 week)
IN NS ns1.example.com.
IN MX 10 smtp.example.com.
;
;
ns1.example.com. IN A 192.168.1.5
smtp.example.com. IN A 192.168.1.5
;
helper.example.com. IN A 192.168.1.5
helper.ocp4.example.com. IN A 192.168.1.5
;
api.ocp4.example.com. IN A 192.168.1.5 1
api-int.ocp4.example.com. IN A 192.168.1.5 2
;
*.apps.ocp4.example.com. IN A 192.168.1.5 3
;
bootstrap.ocp4.example.com. IN A 192.168.1.96 4
;
control-plane0.ocp4.example.com. IN A 192.168.1.97 5
control-plane1.ocp4.example.com. IN A 192.168.1.98 6
control-plane2.ocp4.example.com. IN A 192.168.1.99 7
;
compute0.ocp4.example.com. IN A 192.168.1.11 8
compute1.ocp4.example.com. IN A 192.168.1.7 9
;
;EOF
```

- 1 Provides name resolution for the Kubernetes API. The record refers to the IP address of the API load balancer.

- 2 Provides name resolution for the Kubernetes API. The record refers to the IP address of the API load balancer and is used for internal cluster communications.
- 3 Provides name resolution for the wildcard routes. The record refers to the IP address of the application ingress load balancer. The application ingress load balancer targets the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.



NOTE

In the example, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

- 4 Provides name resolution for the bootstrap machine.
- 5 6 7 Provides name resolution for the control plane machines.
- 8 9 Provides name resolution for the compute machines.

Example DNS PTR record configuration for a user-provisioned cluster

The following example BIND zone file shows sample PTR records for reverse name resolution in a user-provisioned cluster.

Example 15.5. Sample DNS zone database for reverse records

```
$TTL 1W
@ IN SOA ns1.example.com. root (
  2019070700 ; serial
  3H ; refresh (3 hours)
  30M ; retry (30 minutes)
  2W ; expiry (2 weeks)
  1W ) ; minimum (1 week)
IN NS ns1.example.com.
;
5.1.168.192.in-addr.arpa. IN PTR api.ocp4.example.com. 1
5.1.168.192.in-addr.arpa. IN PTR api-int.ocp4.example.com. 2
;
96.1.168.192.in-addr.arpa. IN PTR bootstrap.ocp4.example.com. 3
;
97.1.168.192.in-addr.arpa. IN PTR control-plane0.ocp4.example.com. 4
98.1.168.192.in-addr.arpa. IN PTR control-plane1.ocp4.example.com. 5
99.1.168.192.in-addr.arpa. IN PTR control-plane2.ocp4.example.com. 6
;
11.1.168.192.in-addr.arpa. IN PTR compute0.ocp4.example.com. 7
7.1.168.192.in-addr.arpa. IN PTR compute1.ocp4.example.com. 8
;
;EOF
```

- 1 Provides reverse DNS resolution for the Kubernetes API. The PTR record refers to the record name of the API load balancer.
- 2 Provides reverse DNS resolution for the Kubernetes API. The PTR record refers to the record name of the API load balancer and is used for internal cluster communications.
- 3 Provides reverse DNS resolution for the bootstrap machine.
- 4 5 6 Provides reverse DNS resolution for the control plane machines.
- 7 8 Provides reverse DNS resolution for the compute machines.


**NOTE**

A PTR record is not required for the OpenShift Container Platform application wildcard.

- [Validating DNS resolution for user-provisioned infrastructure](#)

15.3.3.6. Load balancing requirements for user-provisioned infrastructure


Before you install OpenShift Container Platform, you must provision the API and application Ingress load balancing infrastructure. In production scenarios, you can deploy the API and application Ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

**NOTE**

If you want to deploy the API and application Ingress load balancers with a Red Hat Enterprise Linux (RHEL) instance, you must purchase the RHEL subscription separately.

The load balancing infrastructure must meet the following requirements:

1. **API load balancer.** Provides a common endpoint for users, both human and machine, to interact with and configure the platform. Configure the following conditions:
 - Layer 4 load balancing only. This can be referred to as Raw TCP or SSL Passthrough mode.
 - A stateless load balancing algorithm. The options vary based on the load balancer implementation.

**IMPORTANT**

Do not configure session persistence for an API load balancer. Configuring session persistence for a Kubernetes API server might cause performance issues from excess application traffic for your OpenShift Container Platform cluster and the Kubernetes API that runs inside the cluster.

Configure the following ports on both the front and back of the load balancers:

Table 15.17. API load balancer

Port	Back-end machines (pool members)	Internal	External	Description
6443	Bootstrap and control plane. You remove the bootstrap machine from the load balancer after the bootstrap machine initializes the cluster control plane. You must configure the /readyz endpoint for the API server health check probe.	X	X	Kubernetes API server
22623	Bootstrap and control plane. You remove the bootstrap machine from the load balancer after the bootstrap machine initializes the cluster control plane.	X		Machine config server

**NOTE**

The load balancer must be configured to take a maximum of 30 seconds from the time the API server turns off the **/readyz** endpoint to the removal of the API server instance from the pool. Within the time frame after **/readyz** returns an error or becomes healthy, the endpoint must have been removed or added. Probing every 5 or 10 seconds, with two successful requests to become healthy and three to become unhealthy, are well-tested values.

2. **Application Ingress load balancer.** Provides an ingress point for application traffic flowing in from outside the cluster. A working configuration for the Ingress router is required for an OpenShift Container Platform cluster.

Configure the following conditions:

- Layer 4 load balancing only. This can be referred to as Raw TCP or SSL Passthrough mode.
- A connection-based or session-based persistence is recommended, based on the options available and types of applications that will be hosted on the platform.

TIP

If the true IP address of the client can be seen by the application Ingress load balancer, enabling source IP-based session persistence can improve performance for applications that use end-to-end TLS encryption.

Configure the following ports on both the front and back of the load balancers:

Table 15.18. Application Ingress load balancer

Port	Back-end machines (pool members)	Internal	External	Description
443	The machines that run the Ingress Controller pods, compute, or worker, by default.	X	X	HTTPS traffic

Port	Back-end machines (pool members)	Internal	External	Description
80	The machines that run the Ingress Controller pods, compute, or worker, by default.	X	X	HTTP traffic

**NOTE**

If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application Ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes.

15.3.3.6.1. Example load balancer configuration for user-provisioned clusters

This section provides an example API and application Ingress load balancer configuration that meets the load balancing requirements for user-provisioned clusters. The sample is an `/etc/haproxy/haproxy.cfg` configuration for an HAProxy load balancer. The example is not meant to provide advice for choosing one load balancing solution over another.

In the example, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

**NOTE**

If you are using HAProxy as a load balancer and SELinux is set to **enforcing**, you must ensure that the HAProxy service can bind to the configured TCP port by running **setsebool -P haproxy_connect_any=1**.

Example 15.6. Sample API and application Ingress load balancer configuration

```
global
  log      127.0.0.1 local2
  pidfile  /var/run/haproxy.pid
  maxconn  4000
  daemon
defaults
  mode          http
  log           global
  option        dontlognull
  option http-server-close
  option        redispatch
  retries       3
  timeout http-request  10s
  timeout queue        1m
  timeout connect     10s
  timeout client       1m
  timeout server       1m
  timeout http-keep-alive 10s
  timeout check        10s
  maxconn             3000
```

```

listen api-server-6443 1
  bind *:6443
  mode tcp
  option httpchk GET /readyz HTTP/1.0
  option log-health-checks
  balance roundrobin
  server bootstrap bootstrap.ocp4.example.com:6443 verify none check check-ssl inter 10s fall 2
  rise 3 backup 2
  server master0 master0.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
  fall 2 rise 3
  server master1 master1.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
  fall 2 rise 3
  server master2 master2.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
  fall 2 rise 3
listen machine-config-server-22623 3
  bind *:22623
  mode tcp
  server bootstrap bootstrap.ocp4.example.com:22623 check inter 1s backup 4
  server master0 master0.ocp4.example.com:22623 check inter 1s
  server master1 master1.ocp4.example.com:22623 check inter 1s
  server master2 master2.ocp4.example.com:22623 check inter 1s
listen ingress-router-443 5
  bind *:443
  mode tcp
  balance source
  server compute0 compute0.ocp4.example.com:443 check inter 1s
  server compute1 compute1.ocp4.example.com:443 check inter 1s
listen ingress-router-80 6
  bind *:80
  mode tcp
  balance source
  server compute0 compute0.ocp4.example.com:80 check inter 1s
  server compute1 compute1.ocp4.example.com:80 check inter 1s

```

- 1** Port **6443** handles the Kubernetes API traffic and points to the control plane machines.
- 2** **4** The bootstrap entries must be in place before the OpenShift Container Platform cluster installation and they must be removed after the bootstrap process is complete.
- 3** Port **22623** handles the machine config server traffic and points to the control plane machines.
- 5** Port **443** handles the HTTPS traffic and points to the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.
- 6** Port **80** handles the HTTP traffic and points to the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.



NOTE

If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application Ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes.

TIP

If you are using HAProxy as a load balancer, you can check that the **haproxy** process is listening on ports **6443**, **22623**, **443**, and **80** by running **netstat -nltp** on the HAProxy node.

15.3.4. Preparing the user-provisioned infrastructure

Before you install OpenShift Container Platform on user-provisioned infrastructure, you must prepare the underlying infrastructure.

This section provides details about the high-level steps required to set up your cluster infrastructure in preparation for an OpenShift Container Platform installation. This includes configuring IP networking and network connectivity for your cluster nodes, enabling the required ports through your firewall, and setting up the required DNS and load balancing infrastructure.

After preparation, your cluster infrastructure must meet the requirements outlined in the *Requirements for a cluster with user-provisioned infrastructure* section.

Prerequisites

- You have reviewed the [OpenShift Container Platform 4.x Tested Integrations](#) page.
- You have reviewed the infrastructure requirements detailed in the *Requirements for a cluster with user-provisioned infrastructure* section.

Procedure

1. If you are using DHCP to provide the IP networking configuration to your cluster nodes, configure your DHCP service.
 - a. Add persistent IP addresses for the nodes to your DHCP server configuration. In your configuration, match the MAC address of the relevant network interface to the intended IP address for each node.
 - b. When you use DHCP to configure IP addressing for the cluster machines, the machines also obtain the DNS server information through DHCP. Define the persistent DNS server address that is used by the cluster nodes through your DHCP server configuration.

**NOTE**


If you are not using a DHCP service, you must provide the IP networking configuration and the address of the DNS server to the nodes at RHCOS install time. These can be passed as boot arguments if you are installing from an ISO image. See the *Installing RHCOS and starting the OpenShift Container Platform bootstrap process* section for more information about static IP provisioning and advanced networking options.

- c. Define the hostnames of your cluster nodes in your DHCP server configuration. See the *Setting the cluster node hostnames through DHCP* section for details about hostname considerations.

**NOTE**

If you are not using a DHCP service, the cluster nodes obtain their hostname through a reverse DNS lookup.

2. Ensure that your network infrastructure provides the required network connectivity between the cluster components. See the *Networking requirements for user-provisioned infrastructure* section for details about the requirements.
3. Configure your firewall to enable the ports required for the OpenShift Container Platform cluster components to communicate. See *Networking requirements for user-provisioned infrastructure* section for details about the ports that are required.

**IMPORTANT**

By default, port **1936** is accessible for an OpenShift Container Platform cluster, because each control plane node needs access to this port.

Avoid using the Ingress load balancer to expose this port, because doing so might result in the exposure of sensitive information, such as statistics and metrics, related to Ingress Controllers.

4. Setup the required DNS infrastructure for your cluster.
 - a. Configure DNS name resolution for the Kubernetes API, the application wildcard, the bootstrap machine, the control plane machines, and the compute machines.
 - b. Configure reverse DNS resolution for the Kubernetes API, the bootstrap machine, the control plane machines, and the compute machines.
See the *User-provisioned DNS requirements* section for more information about the OpenShift Container Platform DNS requirements.
5. Validate your DNS configuration.
 - a. From your installation node, run DNS lookups against the record names of the Kubernetes API, the wildcard routes, and the cluster nodes. Validate that the IP addresses in the responses correspond to the correct components.
 - b. From your installation node, run reverse DNS lookups against the IP addresses of the load balancer and the cluster nodes. Validate that the record names in the responses correspond to the correct components.
See the *Validating DNS resolution for user-provisioned infrastructure* section for detailed DNS validation steps.
6. Provision the required API and application ingress load balancing infrastructure. See the *Load balancing requirements for user-provisioned infrastructure* section for more information about the requirements.

**NOTE**

Some load balancing solutions require the DNS name resolution for the cluster nodes to be in place before the load balancing is initialized.

Additional resources

- [Requirements for a cluster with user-provisioned infrastructure](#)
- [Installing RHCOS and starting the OpenShift Container Platform bootstrap process](#)
- [Setting the cluster node hostnames through DHCP](#)
- [Advanced RHCOS installation configuration](#)
- [Networking requirements for user-provisioned infrastructure](#)
- [User-provisioned DNS requirements](#)
- [Validating DNS resolution for user-provisioned infrastructure](#)
- [Load balancing requirements for user-provisioned infrastructure](#)

15.3.5. Validating DNS resolution for user-provisioned infrastructure

You can validate your DNS configuration before installing OpenShift Container Platform on user-provisioned infrastructure.



IMPORTANT

The validation steps detailed in this section must succeed before you install your cluster.

Prerequisites

- You have configured the required DNS records for your user-provisioned infrastructure.

Procedure

1. From your installation node, run DNS lookups against the record names of the Kubernetes API, the wildcard routes, and the cluster nodes. Validate that the IP addresses contained in the responses correspond to the correct components.
 - a. Perform a lookup against the Kubernetes API record name. Check that the result points to the IP address of the API load balancer:

```
$ dig +noall +answer @<nameserver_ip> api.<cluster_name>.<base_domain> 1
```

- 1** Replace **<nameserver_ip>** with the IP address of the nameserver, **<cluster_name>** with your cluster name, and **<base_domain>** with your base domain name.

Example output

```
api.ocp4.example.com. 604800 IN A 192.168.1.5
```

- b. Perform a lookup against the Kubernetes internal API record name. Check that the result points to the IP address of the API load balancer:

```
$ dig +noall +answer @<nameserver_ip> api-int.<cluster_name>.<base_domain>
```

Example output

```
api-int.ocp4.example.com. 604800 IN A 192.168.1.5
```

- c. Test an example `*.apps.<cluster_name>.<base_domain>` DNS wildcard lookup. All of the application wildcard lookups must resolve to the IP address of the application ingress load balancer:

```
$ dig +noall +answer @<nameserver_ip> random.apps.<cluster_name>.<base_domain>
```

Example output

```
random.apps.ocp4.example.com. 604800 IN A 192.168.1.5
```



NOTE

In the example outputs, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

You can replace **random** with another wildcard value. For example, you can query the route to the OpenShift Container Platform console:

```
$ dig +noall +answer @<nameserver_ip> console-openshift-console.apps.<cluster_name>.<base_domain>
```

Example output

```
console-openshift-console.apps.ocp4.example.com. 604800 IN A 192.168.1.5
```

- d. Run a lookup against the bootstrap DNS record name. Check that the result points to the IP address of the bootstrap node:

```
$ dig +noall +answer @<nameserver_ip> bootstrap.<cluster_name>.<base_domain>
```

Example output

```
bootstrap.ocp4.example.com. 604800 IN A 192.168.1.96
```

- e. Use this method to perform lookups against the DNS record names for the control plane and compute nodes. Check that the results correspond to the IP addresses of each node.
2. From your installation node, run reverse DNS lookups against the IP addresses of the load balancer and the cluster nodes. Validate that the record names contained in the responses correspond to the correct components.
- a. Perform a reverse lookup against the IP address of the API load balancer. Check that the response includes the record names for the Kubernetes API and the Kubernetes internal API:

```
$ dig +noall +answer @<nameserver_ip> -x 192.168.1.5
```

Example output

```
5.1.168.192.in-addr.arpa. 604800 IN PTR api-int.ocp4.example.com. 1
5.1.168.192.in-addr.arpa. 604800 IN PTR api.ocp4.example.com. 2
```

- 1 Provides the record name for the Kubernetes internal API.
- 2 Provides the record name for the Kubernetes API.



NOTE

A PTR record is not required for the OpenShift Container Platform application wildcard. No validation step is needed for reverse DNS resolution against the IP address of the application ingress load balancer.

- b. Perform a reverse lookup against the IP address of the bootstrap node. Check that the result points to the DNS record name of the bootstrap node:

```
$ dig +noall +answer @<nameserver_ip> -x 192.168.1.96
```

Example output

```
96.1.168.192.in-addr.arpa. 604800 IN PTR bootstrap.ocp4.example.com.
```

- c. Use this method to perform reverse lookups against the IP addresses for the control plane and compute nodes. Check that the results correspond to the DNS record names of each node.

Additional resources

- [User-provisioned DNS requirements](#)
- [Load balancing requirements for user-provisioned infrastructure](#)

15.3.6. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.

**IMPORTANT**

Do not skip this procedure in production environments, where disaster recovery and debugging is required.

**NOTE**

You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.

**NOTE**

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.

**NOTE**

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```

**NOTE**

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

Additional resources

- [Verifying node health](#)

15.3.7. Obtaining the installation program


Before you install OpenShift Container Platform, download the installation file on the host you are using for installation.

Prerequisites

- You have a computer that runs Linux or macOS, with 500 MB of local disk space.

Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select your infrastructure provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.

**IMPORTANT**

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.

**IMPORTANT**

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.

4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#) . This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

15.3.8. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.

**IMPORTANT**

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**.

To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:


```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.



NOTE

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.

5. Move the **oc** binary to a directory on your PATH.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

15.3.9. Manually creating the installation configuration file

Installing the cluster requires that you manually create the installation configuration file.


Prerequisites

- You have an SSH public key on your local machine to provide to the installation program. The key will be used for SSH authentication onto your cluster nodes for debugging and disaster recovery.
- You have obtained the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create an installation directory to store your required installation assets in:

```
$ mkdir <installation_directory>
```



IMPORTANT

You must create a directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

2. Customize the sample **install-config.yaml** file template that is provided and save it in the **<installation_directory>**.



NOTE

You must name this configuration file **install-config.yaml**.

3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the next step of the installation process. You must back it up now.

Additional resources

- [Installation configuration parameters for bare metal](#)

15.3.9.1. Sample install-config.yaml file for bare metal

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.

```

apiVersion: v1
baseDomain: example.com 1
compute: 2
- hyperthreading: Enabled 3
  name: worker
  replicas: 0 4
controlPlane: 5
  hyperthreading: Enabled 6
  name: master
  replicas: 3 7
metadata:
  name: test 8
networking:
  clusterNetwork:
  - cidr: 10.128.0.0/14 9
    hostPrefix: 23 10
  networkType: OVNKubernetes 11
  serviceNetwork: 12
  - 172.30.0.0/16
platform:
  none: {} 13
fips: false 14
pullSecret: '{"auths": ...}' 15
sshKey: 'ssh-ed25519 AAAA...' 16

```

- 1 The base domain of the cluster. All DNS records must be sub-domains of this base and include the cluster name.
- 2 5 The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.
- 3 6 Specifies whether to enable or disable simultaneous multithreading (SMT), or hyperthreading. By default, SMT is enabled to increase the performance of the cores in your machines. You can disable it by setting the parameter value to **Disabled**. If you disable SMT, you must disable it in all cluster machines; this includes both control plane and compute machines.



NOTE

Simultaneous multithreading (SMT) is enabled by default. If SMT is not enabled in your BIOS settings, the **hyperthreading** parameter has no effect.

**IMPORTANT**

If you disable **hyperthreading**, whether in the BIOS or in the **install-config.yaml** file, ensure that your capacity planning accounts for the dramatically decreased machine performance.

- 4 You must set this value to **0** when you install OpenShift Container Platform on user-provisioned infrastructure. In installer-provisioned installations, the parameter controls the number of compute machines that the cluster creates and manages for you. In user-provisioned installations, you must manually deploy the compute machines before you finish installing the cluster.

**NOTE**

If you are installing a three-node cluster, do not deploy any compute machines when you install the Red Hat Enterprise Linux CoreOS (RHCOS) machines.

- 7 The number of control plane machines that you add to the cluster. Because the cluster uses these values as the number of etcd endpoints in the cluster, the value must match the number of control plane machines that you deploy.
- 8 The cluster name that you specified in your DNS records.
- 9 A block of IP addresses from which pod IP addresses are allocated. This block must not overlap with existing physical networks. These IP addresses are used for the pod network. If you need to access the pods from an external network, you must configure load balancers and routers to manage the traffic.

**NOTE**

Class E CIDR range is reserved for a future use. To use the Class E CIDR range, you must ensure your networking environment accepts the IP addresses within the Class E CIDR range.

- 10 The subnet prefix length to assign to each individual node. For example, if **hostPrefix** is set to **23**, then each node is assigned a **/23** subnet out of the given **cidr**, which allows for 510 ($2^{(32 - 23)} - 2$) pod IP addresses. If you are required to provide access to nodes from an external network, configure load balancers and routers to manage the traffic.
- 11 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 12 The IP address pool to use for service IP addresses. You can enter only one IP address pool. This block must not overlap with existing physical networks. If you need to access the services from an external network, configure load balancers and routers to manage the traffic.
- 13 You must set the platform to **none**. You cannot provide additional platform configuration variables for your platform.



IMPORTANT

Clusters that are installed with the platform type **none** are unable to use some features, such as managing compute machines with the Machine API. This limitation applies even if the compute machines that are attached to the cluster are installed on a platform that would normally support the feature. This parameter cannot be changed after installation.

- 14 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.



IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 15 The [pull secret from Red Hat OpenShift Cluster Manager](#). This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.
- 16 The SSH public key for the **core** user in Red Hat Enterprise Linux CoreOS (RHCOS).



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

Additional resources

- See [Load balancing requirements for user-provisioned infrastructure](#) for more information on the API and application ingress load balancing requirements.

15.3.10. Network configuration phases

There are two phases prior to OpenShift Container Platform installation where you can customize the network configuration.

Phase 1

You can customize the following network-related fields in the **install-config.yaml** file before you create the manifest files:

- **networking.networkType**

- **networking.clusterNetwork**
- **networking.serviceNetwork**
- **networking.machineNetwork**

For more information on these fields, refer to *Installation configuration parameters*.



NOTE

Set the **networking.machineNetwork** to match the CIDR that the preferred NIC resides in.



IMPORTANT

The CIDR range **172.17.0.0/16** is reserved by libVirt. You cannot use this range or any range that overlaps with this range for any networks in your cluster.

Phase 2

After creating the manifest files by running **openshift-install create manifests**, you can define a customized Cluster Network Operator manifest with only the fields you want to modify. You can use the manifest to specify advanced network configuration.

You cannot override the values specified in phase 1 in the **install-config.yaml** file during phase 2. However, you can further customize the network plugin during phase 2.

15.3.11. Specifying advanced network configuration

You can use advanced network configuration for your network plugin to integrate your cluster into your existing network environment. You can specify advanced network configuration only before you install the cluster.



IMPORTANT

Customizing your network configuration by modifying the OpenShift Container Platform manifest files created by the installation program is not supported. Applying a manifest file that you create, as in the following procedure, is supported.

Prerequisites

- You have created the **install-config.yaml** file and completed any modifications to it.

Procedure

1. Change to the directory that contains the installation program and create the manifests:

```
$ ./openshift-install create manifests --dir <installation_directory> 1
```

- 1** **<installation_directory>** specifies the name of the directory that contains the **install-config.yaml** file for your cluster.

2. Create a stub manifest file for the advanced network configuration that is named **cluster-network-03-config.yml** in the `<installation_directory>/manifests/` directory:

```
apiVersion: operator.openshift.io/v1
kind: Network
metadata:
  name: cluster
spec:
```

3. Specify the advanced network configuration for your cluster in the **cluster-network-03-config.yml** file, such as in the following example:

Enable IPsec for the OVN-Kubernetes network provider

```
apiVersion: operator.openshift.io/v1
kind: Network
metadata:
  name: cluster
spec:
  defaultNetwork:
    ovnKubernetesConfig:
      ipsecConfig:
        mode: Full
```

4. Optional: Back up the **manifests/cluster-network-03-config.yml** file. The installation program consumes the **manifests/** directory when you create the Ignition config files.

15.3.12. Cluster Network Operator configuration

The configuration for the cluster network is specified as part of the Cluster Network Operator (CNO) configuration and stored in a custom resource (CR) object that is named **cluster**. The CR specifies the fields for the **Network** API in the **operator.openshift.io** API group.

The CNO configuration inherits the following fields during cluster installation from the **Network** API in the **Network.config.openshift.io** API group:

clusterNetwork

IP address pools from which pod IP addresses are allocated.

serviceNetwork

IP address pool for services.

defaultNetwork.type

Cluster network plugin. **OVNKubernetes** is the only supported plugin during installation.

You can specify the cluster network plugin configuration for your cluster by setting the fields for the **defaultNetwork** object in the CNO object named **cluster**.

15.3.12.1. Cluster Network Operator configuration object

The fields for the Cluster Network Operator (CNO) are described in the following table:

Table 15.19. Cluster Network Operator configuration object


Field	Type	Description
metadata.name	string	The name of the CNO object. This name is always cluster .
spec.clusterNetwork	array	A list specifying the blocks of IP addresses from which pod IP addresses are allocated and the subnet prefix length assigned to each individual node in the cluster. For example: <pre>spec: clusterNetwork: - cidr: 10.128.0.0/19 hostPrefix: 23 - cidr: 10.128.32.0/19 hostPrefix: 23</pre>
spec.serviceNetwork	array	A block of IP addresses for services. The OpenShift SDN and OVN-Kubernetes network plugins support only a single IP address block for the service network. For example: <pre>spec: serviceNetwork: - 172.30.0.0/14</pre> <p>You can customize this field only in the install-config.yaml file before you create the manifests. The value is read-only in the manifest file.</p>
spec.defaultNetwork	object	Configures the network plugin for the cluster network.
spec.kubeProxyConfig	object	The fields for this object specify the kube-proxy configuration. If you are using the OVN-Kubernetes cluster network plugin, the kube-proxy configuration has no effect.

defaultNetwork object configuration

The values for the **defaultNetwork** object are defined in the following table:

Table 15.20. **defaultNetwork** object

Field	Type	Description
-------	------	-------------

Field	Type	Description
type	string	<p>OVNKubernetes. The Red Hat OpenShift Networking network plugin is selected during installation. This value cannot be changed after cluster installation.</p> <div style="display: flex; align-items: flex-start;">  <div> <p>NOTE</p> <p>OpenShift Container Platform uses the OVN-Kubernetes network plugin by default. OpenShift SDN is no longer available as an installation choice for new clusters.</p> </div> </div>
ovnKubernetesConfig	object	This object is only valid for the OVN-Kubernetes network plugin.

Configuration for the OVN-Kubernetes network plugin

The following table describes the configuration fields for the OVN-Kubernetes network plugin:

Table 15.21. ovnKubernetesConfig object

Field	Type	Description
mtu	integer	<p>The maximum transmission unit (MTU) for the Geneve (Generic Network Virtualization Encapsulation) overlay network. This is detected automatically based on the MTU of the primary network interface. You do not normally need to override the detected MTU.</p> <p>If the auto-detected value is not what you expect it to be, confirm that the MTU on the primary network interface on your nodes is correct. You cannot use this option to change the MTU value of the primary network interface on the nodes.</p> <p>If your cluster requires different MTU values for different nodes, you must set this value to 100 less than the lowest MTU value in your cluster. For example, if some nodes in your cluster have an MTU of 9001, and some have an MTU of 1500, you must set this value to 1400.</p>
genevePort	integer	The port to use for all Geneve packets. The default value is 6081 . This value cannot be changed after cluster installation.


Field	Type	Description
ipsecConfig	object	Specify a configuration object for customizing the IPsec configuration.
ipv4	object	Specifies a configuration object for IPv4 settings.
ipv6	object	Specifies a configuration object for IPv6 settings.
policyAuditConfig	object	Specify a configuration object for customizing network policy audit logging. If unset, the defaults audit log settings are used.
gatewayConfig	object	Optional: Specify a configuration object for customizing how egress traffic is sent to the node gateway.  <p>NOTE</p> <p>While migrating egress traffic, you can expect some disruption to workloads and service traffic until the Cluster Network Operator (CNO) successfully rolls out the changes.</p>

Table 15.22. `ovnKubernetesConfig.ipv4` object

Field	Type	Description
internalTransitSwitchSubnet	string	If your existing network infrastructure overlaps with the 100.88.0.0/16 IPv4 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. The subnet for the distributed transit switch that enables east-west traffic. This subnet cannot overlap with any other subnets used by OVN-Kubernetes or on the host itself. It must be large enough to accommodate one IP address per node in your cluster. The default value is 100.88.0.0/16 .
internalJoinSubnet	string	If your existing network infrastructure overlaps with the 100.64.0.0/16 IPv4 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. You must ensure that the IP address range does not overlap with any other subnet used by your OpenShift Container Platform installation. The IP address range must be larger than the maximum number of nodes that can be added to the cluster. For example, if the clusterNetwork.cidr value is 10.128.0.0/14 and the clusterNetwork.hostPrefix value is /23 , then the maximum number of nodes is 2⁽²³⁻¹⁴⁾=512 . The default value is 100.64.0.0/16 .

Table 15.23. `ovnKubernetesConfig.ipv6` object

Field	Type	Description
internalTransitSwitchSubnet	string	<p>If your existing network infrastructure overlaps with the fd97::/64 IPv6 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. The subnet for the distributed transit switch that enables east-west traffic. This subnet cannot overlap with any other subnets used by OVN-Kubernetes or on the host itself. It must be large enough to accommodate one IP address per node in your cluster.</p> <p>The default value is fd97::/64.</p>
internalJoinSubnet	string	<p>If your existing network infrastructure overlaps with the fd98::/64 IPv6 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. You must ensure that the IP address range does not overlap with any other subnet used by your OpenShift Container Platform installation. The IP address range must be larger than the maximum number of nodes that can be added to the cluster.</p> <p>The default value is fd98::/64.</p>

Table 15.24. `policyAuditConfig` object

Field	Type	Description
rateLimit	integer	The maximum number of messages to generate every second per node. The default value is 20 messages per second.
maxFileSize	integer	The maximum size for the audit log in bytes. The default value is 50000000 or 50 MB.
maxLogFiles	integer	The maximum number of log files that are retained.
destination	string	<p>One of the following additional audit log targets:</p> <p>libc The libc syslog() function of the journald process on the host.</p> <p>udp:<host>:<port> A syslog server. Replace <host>:<port> with the host and port of the syslog server.</p> <p>unix:<file> A Unix Domain Socket file specified by <file>.</p> <p>null Do not send the audit logs to any additional target.</p>
syslogFacility	string	The syslog facility, such as kern , as defined by RFC5424. The default value is local0 .

Table 15.25. gatewayConfig object

Field	Type	Description
routingViaHost	boolean	<p>Set this field to true to send egress traffic from pods to the host networking stack. For highly-specialized installations and applications that rely on manually configured routes in the kernel routing table, you might want to route egress traffic to the host networking stack. By default, egress traffic is processed in OVN to exit the cluster and is not affected by specialized routes in the kernel routing table. The default value is false.</p> <p>This field has an interaction with the Open vSwitch hardware offloading feature. If you set this field to true, you do not receive the performance benefits of the offloading because egress traffic is processed by the host networking stack.</p>
ipForwarding	object	<p>You can control IP forwarding for all traffic on OVN-Kubernetes managed interfaces by using the ipForwarding specification in the Network resource. Specify Restricted to only allow IP forwarding for Kubernetes related traffic. Specify Global to allow forwarding of all IP traffic. For new installations, the default is Restricted. For updates to OpenShift Container Platform 4.14 or later, the default is Global.</p>
ipv4	object	<p>Optional: Specify an object to configure the internal OVN-Kubernetes masquerade address for host to service traffic for IPv4 addresses.</p>
ipv6	object	<p>Optional: Specify an object to configure the internal OVN-Kubernetes masquerade address for host to service traffic for IPv6 addresses.</p>

Table 15.26. gatewayConfig.ipv4 object

Field	Type	Description
internalMasqueradeSubnet	string	<p>The masquerade IPv4 addresses that are used internally to enable host to service traffic. The host is configured with these IP addresses as well as the shared gateway bridge interface. The default value is 169.254.169.0/29.</p>

Table 15.27. gatewayConfig.ipv6 object


Field	Type	Description
internalMasqueradeSubnet	string	<p>The masquerade IPv6 addresses that are used internally to enable host to service traffic. The host is configured with these IP addresses as well as the shared gateway bridge interface. The default value is fd69::/125.</p>

Table 15.28. `ipsecConfig` object

Field	Type	Description
<code>mode</code>	<code>string</code>	<p>Specifies the behavior of the IPsec implementation. Must be one of the following values:</p> <ul style="list-style-type: none"> ● Disabled: IPsec is not enabled on cluster nodes. ● External: IPsec is enabled for network traffic with external hosts. ● Full: IPsec is enabled for pod traffic and network traffic with external hosts.

Example OVN-Kubernetes configuration with IPsec enabled

```
defaultNetwork:
  type: OVNKubernetes
  ovnKubernetesConfig:
    mtu: 1400
    genevePort: 6081
    ipsecConfig:
      mode: Full
```




IMPORTANT

Using OVNKubernetes can lead to a stack exhaustion problem on IBM Power®.

kubeProxyConfig object configuration (OpenShiftSDN container network interface only)

The values for the `kubeProxyConfig` object are defined in the following table:

Table 15.29. `kubeProxyConfig` object

Field	Type	Description
<code>iptablesSyncPeriod</code>	<code>string</code>	<p>The refresh period for <code>iptables</code> rules. The default value is 30s. Valid suffixes include s, m, and h and are described in the Go time package documentation.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p>NOTE</p> <p>Because of performance improvements introduced in OpenShift Container Platform 4.3 and greater, adjusting the <code>iptablesSyncPeriod</code> parameter is no longer necessary.</p> </div> </div>

Field	Type	Description
proxyArguments.iptables-min-sync-period	array	<p>The minimum duration before refreshing iptables rules. This field ensures that the refresh does not happen too frequently. Valid suffixes include s, m, and h and are described in the Go time package. The default value is:</p> <pre>kubeProxyConfig: proxyArguments: iptables-min-sync-period: - 0s</pre>

15.3.13. Creating the Ignition config files

Because you must manually start the cluster machines, you must generate the Ignition config files that the cluster needs to make its machines.



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

Prerequisites


- Obtain the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

- Obtain the Ignition config files:

```
$ ./openshift-install create ignition-configs --dir <installation_directory> 1
```


- 1** For **<installation_directory>**, specify the directory name to store the files that the installation program creates.



IMPORTANT

If you created an **install-config.yaml** file, specify the directory that contains it. Otherwise, specify an empty directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

The following files are generated in the directory:



15.3.14. Installing RHCOS and starting the OpenShift Container Platform bootstrap process

To install OpenShift Container Platform on bare metal infrastructure that you provision, you must install Red Hat Enterprise Linux CoreOS (RHCOS) on the machines. When you install RHCOS, you must provide the Ignition config file that was generated by the OpenShift Container Platform installation program for the type of machine you are installing. If you have configured suitable networking, DNS, and load balancing infrastructure, the OpenShift Container Platform bootstrap process begins automatically after the RHCOS machines have rebooted.

To install RHCOS on the machines, follow either the steps to use an ISO image or network PXE booting.



NOTE

The compute node deployment steps included in this installation document are RHCOS-specific. If you choose instead to deploy RHEL-based compute nodes, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Only RHEL 8 compute machines are supported.

You can configure RHCOS during ISO and PXE installations by using the following methods:

- **Kernel arguments:** You can use kernel arguments to provide installation-specific information. For example, you can specify the locations of the RHCOS installation files that you uploaded to your HTTP server and the location of the Ignition config file for the type of node you are installing. For a PXE installation, you can use the **APPEND** parameter to pass the arguments to the kernel of the live installer. For an ISO installation, you can interrupt the live installation boot process to add the kernel arguments. In both installation cases, you can use special **coreos.inst.*** arguments to direct the live installer, as well as standard installation boot arguments for turning standard kernel services on or off.
- **Ignition configs:** OpenShift Container Platform Ignition config files (***.ign**) are specific to the type of node you are installing. You pass the location of a bootstrap, control plane, or compute

node Ignition config file during the RHCOS installation so that it takes effect on first boot. In special cases, you can create a separate, limited Ignition config to pass to the live system. That Ignition config could do a certain set of tasks, such as reporting success to a provisioning system after completing installation. This special Ignition config is consumed by the **coreos-installer** to be applied on first boot of the installed system. Do not provide the standard control plane and compute node Ignition configs to the live ISO directly.

- **coreos-installer**: You can boot the live ISO installer to a shell prompt, which allows you to prepare the permanent system in a variety of ways before first boot. In particular, you can run the **coreos-installer** command to identify various artifacts to include, work with disk partitions, and set up networking. In some cases, you can configure features on the live system and copy them to the installed system.

Whether to use an ISO or PXE install depends on your situation. A PXE install requires an available DHCP service and more preparation, but can make the installation process more automated. An ISO install is a more manual process and can be inconvenient if you are setting up more than a few machines.



NOTE

As of OpenShift Container Platform 4.6, the RHCOS ISO and other installation artifacts provide support for installation on disks with 4K sectors.

15.3.14.1. Installing RHCOS by using an ISO image

You can use an ISO image to install RHCOS on the machines.

Prerequisites

- You have created the Ignition config files for your cluster.
- You have configured suitable network, DNS and load balancing infrastructure.
- You have an HTTP server that can be accessed from your computer, and from the machines that you create.
- You have reviewed the *Advanced RHCOS installation configuration* section for different ways to configure features, such as networking and disk partitioning.

Procedure

1. Obtain the SHA512 digest for each of your Ignition config files. For example, you can use the following on a system running Linux to get the SHA512 digest for your **bootstrap.ign** Ignition config file:

```
$ sha512sum <installation_directory>/bootstrap.ign
```

The digests are provided to the **coreos-installer** in a later step to validate the authenticity of the Ignition config files on the cluster nodes.

2. Upload the bootstrap, control plane, and compute node Ignition config files that the installation program created to your HTTP server. Note the URLs of these files.



IMPORTANT

You can add or change configuration settings in your Ignition configs before saving them to your HTTP server. If you plan to add more compute machines to your cluster after you finish installation, do not delete these files.

- From the installation host, validate that the Ignition config files are available on the URLs. The following example gets the Ignition config file for the bootstrap node:

```
$ curl -k http://<HTTP_server>/bootstrap.ign 1
```

Example output

```
% Total    % Received % Xferd Average Speed   Time    Time     Time  Current
           Dload  Upload  Total   Spent    Left    Speed
  0   0   0   0   0   0   0   0   0  --:--:-- --:--:-- --:--:--    0{"ignition":
{"version":"3.2.0"},"passwd":{"users":[{"name":"core","sshAuthorizedKeys":["ssh-rsa...
```


Replace **bootstrap.ign** with **master.ign** or **worker.ign** in the command to validate that the Ignition config files for the control plane and compute nodes are also available.

- Although it is possible to obtain the RHCOS images that are required for your preferred method of installing operating system instances from the [RHCOS image mirror](#) page, the recommended way to obtain the correct version of your RHCOS images are from the output of **openshift-install** command:

```
$ openshift-install coreos print-stream-json | grep '\.iso[^\.]'
```

Example output

```
"location": "<url>/art/storage/releases/rhcos-4.16-aarch64/<release>/aarch64/rhcos-
<release>-live.aarch64.iso",
"location": "<url>/art/storage/releases/rhcos-4.16-ppc64le/<release>/ppc64le/rhcos-
<release>-live.ppc64le.iso",
"location": "<url>/art/storage/releases/rhcos-4.16-s390x/<release>/s390x/rhcos-<release>-
live.s390x.iso",
"location": "<url>/art/storage/releases/rhcos-4.16/<release>/x86_64/rhcos-<release>-
live.x86_64.iso",
```



IMPORTANT

The RHCOS images might not change with every release of OpenShift Container Platform. You must download images with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Use the image versions that match your OpenShift Container Platform version if they are available. Use only ISO images for this procedure. RHCOS qcow2 images are not supported for this installation type.

ISO file names resemble the following example:

rhcos-<version>-live.<architecture>.iso

- Use the ISO to start the RHCOS installation. Use one of the following installation options:

- Burn the ISO image to a disk and boot it directly.
 - Use ISO redirection by using a lights-out management (LOM) interface.
6. Boot the RHCOS ISO image without specifying any options or interrupting the live boot sequence. Wait for the installer to boot into a shell prompt in the RHCOS live environment.



NOTE

It is possible to interrupt the RHCOS installation boot process to add kernel arguments. However, for this ISO procedure you should use the **coreos-installer** command as outlined in the following steps, instead of adding kernel arguments.

7. Run the **coreos-installer** command and specify the options that meet your installation requirements. At a minimum, you must specify the URL that points to the Ignition config file for the node type, and the device that you are installing to:

```
$ sudo coreos-installer install --ignition-url=http://<HTTP_server>/<node_type>.ign <device>
--ignition-hash=sha512-<digest> 1 2
```

- 1 1 You must run the **coreos-installer** command by using **sudo**, because the **core** user does not have the required root privileges to perform the installation.

- 2 The **--ignition-hash** option is required when the Ignition config file is obtained through an HTTP URL to validate the authenticity of the Ignition config file on the cluster node. **<digest>** is the Ignition config file SHA512 digest obtained in a preceding step.



NOTE

If you want to provide your Ignition config files through an HTTPS server that uses TLS, you can add the internal certificate authority (CA) to the system trust store before running **coreos-installer**.

The following example initializes a bootstrap node installation to the **/dev/sda** device. The Ignition config file for the bootstrap node is obtained from an HTTP web server with the IP address 192.168.1.2:

```
$ sudo coreos-installer install --ignition-
url=http://192.168.1.2:80/installation_directory/bootstrap.ign /dev/sda --ignition-hash=sha512-
a5a2d43879223273c9b60af66b44202a1d1248fc01cf156c46d4a79f552b6bad47bc8cc78ddf011
6e80c59d2ea9e32ba53bc807afbca581aa059311def2c3e3b
```

8. Monitor the progress of the RHCOS installation on the console of the machine.



IMPORTANT


Be sure that the installation is successful on each node before commencing with the OpenShift Container Platform installation. Observing the installation process can also help to determine the cause of RHCOS installation issues that might arise.

9. After RHCOS installs, you must reboot the system. During the system reboot, it applies the Ignition config file that you specified.
10. Check the console output to verify that Ignition ran.

Example command

```
Ignition: ran on 2022/03/14 14:48:33 UTC (this boot)
Ignition: user-provided config was applied
```

11. Continue to create the other machines for your cluster.



IMPORTANT

You must create the bootstrap and control plane machines at this time. If the control plane machines are not made schedulable, also create at least two compute machines before you install OpenShift Container Platform.

If the required network, DNS, and load balancer infrastructure are in place, the OpenShift Container Platform bootstrap process begins automatically after the RHCOS nodes have rebooted.



NOTE

RHCOS nodes do not include a default password for the **core** user. You can access the nodes by running **ssh core@<node>.<cluster_name>.<base_domain>** as a user with access to the SSH private key that is paired to the public key that you specified in your **install_config.yaml** file. OpenShift Container Platform 4 cluster nodes running RHCOS are immutable and rely on Operators to apply cluster changes. Accessing cluster nodes by using SSH is not recommended. However, when investigating installation issues, if the OpenShift Container Platform API is not available, or the kubelet is not properly functioning on a target node, SSH access might be required for debugging or disaster recovery.

15.3.14.2. Installing RHCOS by using PXE or iPXE booting

You can use PXE or iPXE booting to install RHCOS on the machines.

Prerequisites

- You have created the Ignition config files for your cluster.
- You have configured suitable network, DNS and load balancing infrastructure.
- You have configured suitable PXE or iPXE infrastructure.
- You have an HTTP server that can be accessed from your computer, and from the machines that you create.
- You have reviewed the *Advanced RHCOS installation configuration* section for different ways to configure features, such as networking and disk partitioning.

Procedure

1. Upload the bootstrap, control plane, and compute node Ignition config files that the installation program created to your HTTP server. Note the URLs of these files.



IMPORTANT

You can add or change configuration settings in your Ignition configs before saving them to your HTTP server. If you plan to add more compute machines to your cluster after you finish installation, do not delete these files.

2. From the installation host, validate that the Ignition config files are available on the URLs. The following example gets the Ignition config file for the bootstrap node:

```
$ curl -k http://<HTTP_server>/bootstrap.ign 1
```

Example output

```
% Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
           Dload  Upload  Total      Spent    Left  Speed
  0   0   0   0   0   0   0   0  --:--:--  --:--:--  --:--:--    0{"ignition":
{"version":"3.2.0"},"passwd":{"users":[{"name":"core","sshAuthorizedKeys":["ssh-rsa...
```

Replace **bootstrap.ign** with **master.ign** or **worker.ign** in the command to validate that the Ignition config files for the control plane and compute nodes are also available.

3. Although it is possible to obtain the RHCOS **kernel**, **initramfs** and **rootfs** files that are required for your preferred method of installing operating system instances from the [RHCOS image mirror](#) page, the recommended way to obtain the correct version of your RHCOS files are from the output of **openshift-install** command:

```
$ openshift-install coreos print-stream-json | grep -Eo "https.*(kernel-|initramfs-|rootfs-)\w+(\.img)?"
```

Example output

```
"<url>/art/storage/releases/rhcos-4.16-aarch64/<release>/aarch64/rhcos-<release>-live-
kernel-aarch64"
"<url>/art/storage/releases/rhcos-4.16-aarch64/<release>/aarch64/rhcos-<release>-live-
initramfs.aarch64.img"
"<url>/art/storage/releases/rhcos-4.16-aarch64/<release>/aarch64/rhcos-<release>-live-
rootfs.aarch64.img"
"<url>/art/storage/releases/rhcos-4.16-ppc64le/49.84.202110081256-0/ppc64le/rhcos-
<release>-live-kernel-ppc64le"
"<url>/art/storage/releases/rhcos-4.16-ppc64le/<release>/ppc64le/rhcos-<release>-live-
initramfs.ppc64le.img"
"<url>/art/storage/releases/rhcos-4.16-ppc64le/<release>/ppc64le/rhcos-<release>-live-
rootfs.ppc64le.img"
"<url>/art/storage/releases/rhcos-4.16-s390x/<release>/s390x/rhcos-<release>-live-kernel-
s390x"
"<url>/art/storage/releases/rhcos-4.16-s390x/<release>/s390x/rhcos-<release>-live-
initramfs.s390x.img"
"<url>/art/storage/releases/rhcos-4.16-s390x/<release>/s390x/rhcos-<release>-live-
rootfs.s390x.img"
"<url>/art/storage/releases/rhcos-4.16/<release>/x86_64/rhcos-<release>-live-kernel-
```

```
x86_64"
"<url>/art/storage/releases/rhcos-4.16/<release>/x86_64/rhcos-<release>-live-
initramfs.x86_64.img"
"<url>/art/storage/releases/rhcos-4.16/<release>/x86_64/rhcos-<release>-live-
rootfs.x86_64.img"
```




IMPORTANT

The RHCOS artifacts might not change with every release of OpenShift Container Platform. You must download images with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Only use the appropriate **kernel**, **initramfs**, and **rootfs** artifacts described below for this procedure. RHCOS QCOW2 images are not supported for this installation type.

The file names contain the OpenShift Container Platform version number. They resemble the following examples:

- **kernel:** `rhcos-<version>-live-kernel-<architecture>`
- **initramfs:** `rhcos-<version>-live-initramfs.<architecture>.img`
- **rootfs:** `rhcos-<version>-live-rootfs.<architecture>.img`

4. Upload the **rootfs**, **kernel**, and **initramfs** files to your HTTP server.



IMPORTANT

If you plan to add more compute machines to your cluster after you finish installation, do not delete these files.

5. Configure the network boot infrastructure so that the machines boot from their local disks after RHCOS is installed on them.
6. Configure PXE or iPXE installation for the RHCOS images and begin the installation. Modify one of the following example menu entries for your environment and verify that the image and Ignition files are properly accessible:

- For PXE (**x86_64**):

```
DEFAULT pxeboot
TIMEOUT 20
PROMPT 0
LABEL pxeboot
  KERNEL http://<HTTP_server>/rhcos-<version>-live-kernel-<architecture> 1
  APPEND initrd=http://<HTTP_server>/rhcos-<version>-live-initramfs.
<architecture>.img coreos.live.rootfs_url=http://<HTTP_server>/rhcos-<version>-live-
rootfs.<architecture>.img coreos.inst.install_dev=/dev/sda
coreos.inst.ignition_url=http://<HTTP_server>/bootstrap.ign 2 3
```

- 1** Specify the location of the live **kernel** file that you uploaded to your HTTP server. The URL must be HTTP, TFTP, or FTP; HTTPS and NFS are not supported.
- 2** If you use multiple NICs, specify a single interface in the **ip** option. For example, to use DHCP on a NIC that is named `eno1`, set `ip=eno1:dhcp`

DHCP on a NIC that is named **eno1**, set **ip=eno1:dhcp**.

- 3 Specify the locations of the RHCOS files that you uploaded to your HTTP server. The **initrd** parameter value is the location of the **initramfs** file, the **coreos.live.rootfs_url** parameter value is the location of the **rootfs** file, and the **coreos.inst.ignition_url** parameter value is the location of the bootstrap Ignition config file. You can also add more kernel arguments to the **APPEND** line to configure networking or other boot options.



NOTE

This configuration does not enable serial console access on machines with a graphical console. To configure a different console, add one or more **console=** arguments to the **APPEND** line. For example, add **console=tty0 console=ttyS0** to set the first PC serial port as the primary console and the graphical console as a secondary console. For more information, see [How does one set up a serial terminal and/or console in Red Hat Enterprise Linux?](#) and "Enabling the serial console for PXE and ISO installation" in the "Advanced RHCOS installation configuration" section.

- For iPXE (**x86_64 + aarch64**):

```
kernel http://<HTTP_server>/rhcos-<version>-live-kernel-<architecture> initrd=main
coreos.live.rootfs_url=http://<HTTP_server>/rhcos-<version>-live-rootfs.
<architecture>.img coreos.inst.install_dev=/dev/sda
coreos.inst.ignition_url=http://<HTTP_server>/bootstrap.ign 1 2
initrd --name main http://<HTTP_server>/rhcos-<version>-live-initramfs.
<architecture>.img 3
boot
```

- 1 Specify the locations of the RHCOS files that you uploaded to your HTTP server. The **kernel** parameter value is the location of the **kernel** file, the **initrd=main** argument is needed for booting on UEFI systems, the **coreos.live.rootfs_url** parameter value is the location of the **rootfs** file, and the **coreos.inst.ignition_url** parameter value is the location of the bootstrap Ignition config file.
- 2 If you use multiple NICs, specify a single interface in the **ip** option. For example, to use DHCP on a NIC that is named **eno1**, set **ip=eno1:dhcp**.
- 3 Specify the location of the **initramfs** file that you uploaded to your HTTP server.



NOTE

This configuration does not enable serial console access on machines with a graphical console. To configure a different console, add one or more **console=** arguments to the **kernel** line. For example, add **console=tty0 console=ttyS0** to set the first PC serial port as the primary console and the graphical console as a secondary console. For more information, see [How does one set up a serial terminal and/or console in Red Hat Enterprise Linux?](#) and "Enabling the serial console for PXE and ISO installation" in the "Advanced RHCOS installation configuration" section.

**NOTE**

To network boot the CoreOS **kernel** on **aarch64** architecture, you need to use a version of iPXE build with the **IMAGE_GZIP** option enabled. See [IMAGE_GZIP option in iPXE](#).

- For PXE (with UEFI and Grub as second stage) on **aarch64**:

```

menuentry 'Install CoreOS' {
  linux rhcos-<version>-live-kernel-<architecture>
  coreos.live.rootfs_url=http://<HTTP_server>/rhcos-<version>-live-rootfs.
  <architecture>.img coreos.inst.install_dev=/dev/sda
  coreos.inst.ignition_url=http://<HTTP_server>/bootstrap.ign 1 2
  initrd rhcos-<version>-live-initramfs.<architecture>.img 3
}

```

- 1** Specify the locations of the RHCOS files that you uploaded to your HTTP/TFTP server. The **kernel** parameter value is the location of the **kernel** file on your TFTP server. The **coreos.live.rootfs_url** parameter value is the location of the **rootfs** file, and the **coreos.inst.ignition_url** parameter value is the location of the bootstrap Ignition config file on your HTTP Server.
- 2** If you use multiple NICs, specify a single interface in the **ip** option. For example, to use DHCP on a NIC that is named **eno1**, set **ip=eno1:dhcp**.
- 3** Specify the location of the **initramfs** file that you uploaded to your TFTP server.

7. Monitor the progress of the RHCOS installation on the console of the machine.

**IMPORTANT**

Be sure that the installation is successful on each node before commencing with the OpenShift Container Platform installation. Observing the installation process can also help to determine the cause of RHCOS installation issues that might arise.

8. After RHCOS installs, the system reboots. During reboot, the system applies the Ignition config file that you specified.
9. Check the console output to verify that Ignition ran.


Example command

```

Ignition: ran on 2022/03/14 14:48:33 UTC (this boot)
Ignition: user-provided config was applied

```

10. Continue to create the machines for your cluster.

**IMPORTANT**

You must create the bootstrap and control plane machines at this time. If the control plane machines are not made schedulable, also create at least two compute machines before you install the cluster.

If the required network, DNS, and load balancer infrastructure are in place, the OpenShift Container Platform bootstrap process begins automatically after the RHCOS nodes have rebooted.



NOTE

RHCOS nodes do not include a default password for the **core** user. You can access the nodes by running **ssh core@<node>.<cluster_name>.<base_domain>** as a user with access to the SSH private key that is paired to the public key that you specified in your **install_config.yaml** file. OpenShift Container Platform 4 cluster nodes running RHCOS are immutable and rely on Operators to apply cluster changes. Accessing cluster nodes by using SSH is not recommended. However, when investigating installation issues, if the OpenShift Container Platform API is not available, or the kubelet is not properly functioning on a target node, SSH access might be required for debugging or disaster recovery.

15.3.14.3. Advanced RHCOS installation configuration

A key benefit for manually provisioning the Red Hat Enterprise Linux CoreOS (RHCOS) nodes for OpenShift Container Platform is to be able to do configuration that is not available through default OpenShift Container Platform installation methods. This section describes some of the configurations that you can do using techniques that include:

- Passing kernel arguments to the live installer
- Running **coreos-installer** manually from the live system
- Customizing a live ISO or PXE boot image

The advanced configuration topics for manual Red Hat Enterprise Linux CoreOS (RHCOS) installations detailed in this section relate to disk partitioning, networking, and using Ignition configs in different ways.

15.3.14.3.1. Using advanced networking options for PXE and ISO installations

Networking for OpenShift Container Platform nodes uses DHCP by default to gather all necessary configuration settings. To set up static IP addresses or configure special settings, such as bonding, you can do one of the following:

- Pass special kernel parameters when you boot the live installer.
- Use a machine config to copy networking files to the installed system.
- Configure networking from a live installer shell prompt, then copy those settings to the installed system so that they take effect when the installed system first boots.

To configure a PXE or iPXE installation, use one of the following options:

- See the "Advanced RHCOS installation reference" tables.
- Use a machine config to copy networking files to the installed system.

To configure an ISO installation, use the following procedure.

Procedure

1. Boot the ISO installer.
2. From the live system shell prompt, configure networking for the live system using available RHEL tools, such as **nmcli** or **nmtui**.
3. Run the **coreos-installer** command to install the system, adding the **--copy-network** option to copy networking configuration. For example:

```
$ sudo coreos-installer install --copy-network \
  --ignition-url=http://host/worker.ign /dev/disk/by-id/scsi-<serial_number>
```



IMPORTANT

The **--copy-network** option only copies networking configuration found under **/etc/NetworkManager/system-connections**. In particular, it does not copy the system hostname.


4. Reboot into the installed system.

Additional resources

- See [Getting started with nmcli](#) and [Getting started with nmtui](#) in the RHEL 8 documentation for more information about the **nmcli** and **nmtui** tools.

15.3.14.3.2. Disk partitioning

Disk partitions are created on OpenShift Container Platform cluster nodes during the Red Hat Enterprise Linux CoreOS (RHCOS) installation. Each RHCOS node of a particular architecture uses the same partition layout, unless you override the default partitioning configuration. During the RHCOS installation, the size of the root file system is increased to use any remaining available space on the target device.



IMPORTANT


The use of a custom partition scheme on your node might result in OpenShift Container Platform not monitoring or alerting on some node partitions. If you override the default partitioning, see [Understanding OpenShift File System Monitoring \(eviction conditions\)](#) for more information about how OpenShift Container Platform monitors your host file systems.

OpenShift Container Platform monitors the following two filesystem identifiers:

- **nodefs**, which is the filesystem that contains **/var/lib/kubelet**
- **imagefs**, which is the filesystem that contains **/var/lib/containers**

For the default partition scheme, **nodefs** and **imagefs** monitor the same root filesystem, **/**.

To override the default partitioning when installing RHCOS on an OpenShift Container Platform cluster node, you must create separate partitions. Consider a situation where you want to add a separate storage partition for your containers and container images. For example, by mounting **/var/lib/containers** in a separate partition, the kubelet separately monitors **/var/lib/containers** as the **imagefs** directory and the root file system as the **nodefs** directory.

**IMPORTANT**


If you have resized your disk size to host a larger file system, consider creating a separate **/var/lib/containers** partition. Consider resizing a disk that has an **xfs** format to reduce CPU time issues caused by a high number of allocation groups.

15.3.14.3.2.1. Creating a separate /var partition

In general, you should use the default disk partitioning that is created during the RHCOS installation. However, there are cases where you might want to create a separate partition for a directory that you expect to grow.

OpenShift Container Platform supports the addition of a single partition to attach storage to either the **/var** directory or a subdirectory of **/var**. For example:

- **/var/lib/containers**: Holds container-related content that can grow as more images and containers are added to a system.
- **/var/lib/etcd**: Holds data that you might want to keep separate for purposes such as performance optimization of etcd storage.
- **/var**: Holds data that you might want to keep separate for purposes such as auditing.

**IMPORTANT**

For disk sizes larger than 100GB, and especially larger than 1TB, create a separate **/var** partition.

Storing the contents of a **/var** directory separately makes it easier to grow storage for those areas as needed and reinstall OpenShift Container Platform at a later date and keep that data intact. With this method, you will not have to pull all your containers again, nor will you have to copy massive log files when you update systems.

The use of a separate partition for the **/var** directory or a subdirectory of **/var** also prevents data growth in the partitioned directory from filling up the root file system.

The following procedure sets up a separate **/var** partition by adding a machine config manifest that is wrapped into the Ignition config file for a node type during the preparation phase of an installation.

Procedure

1. On your installation host, change to the directory that contains the OpenShift Container Platform installation program and generate the Kubernetes manifests for the cluster:

```
$ openshift-install create manifests --dir <installation_directory>
```

2. Create a Butane config that configures the additional partition. For example, name the file **\$HOME/clusterconfig/98-var-partition.bu**, change the disk device name to the name of the storage device on the **worker** systems, and set the storage size as appropriate. This example places the **/var** directory on a separate partition:

```
variant: openshift
version: 4.16.0
metadata:
  labels:
```




```

machineconfiguration.openshift.io/role: worker
name: 98-var-partition
storage:
  disks:
    - device: /dev/disk/by-id/<device_name> ❶
      partitions:
        - label: var
          start_mib: <partition_start_offset> ❷
          size_mib: <partition_size> ❸
          number: 5
      filesystems:
        - device: /dev/disk/by-partlabel/var
          path: /var
          format: xfs
          mount_options: [defaults, prjquota] ❹
          with_mount_unit: true

```

- ❶ The storage device name of the disk that you want to partition.
- ❷ When adding a data partition to the boot disk, a minimum offset value of 25000 mebibytes is recommended. The root file system is automatically resized to fill all available space up to the specified offset. If no offset value is specified, or if the specified value is smaller than the recommended minimum, the resulting root file system will be too small, and future reinstalls of RHCOS might overwrite the beginning of the data partition.
- ❸ The size of the data partition in mebibytes.
- ❹ The **prjquota** mount option must be enabled for filesystems used for container storage.



NOTE

When creating a separate **/var** partition, you cannot use different instance types for compute nodes, if the different instance types do not have the same device name.

3. Create a manifest from the Butane config and save it to the **clusterconfig/openshift** directory. For example, run the following command:

```
$ butane $HOME/clusterconfig/98-var-partition.bu -o $HOME/clusterconfig/openshift/98-var-partition.yaml
```


4. Create the Ignition config files:

```
$ openshift-install create ignition-configs --dir <installation_directory> ❶
```

- ❶ For **<installation_directory>**, specify the same installation directory.

Ignition config files are created for the bootstrap, control plane, and compute nodes in the installation directory:

```
├── auth
```



The files in the `<installation_directory>/manifest` and `<installation_directory>/openshift` directories are wrapped into the Ignition config files, including the file that contains the **98-var-partition** custom **MachineConfig** object.

Next steps

- You can apply the custom disk partitioning by referencing the Ignition config files during the RHCOS installations.

15.3.14.3.2.2. Retaining existing partitions

For an ISO installation, you can add options to the **coreos-installer** command that cause the installer to maintain one or more existing partitions. For a PXE installation, you can add **coreos.inst.*** options to the **APPEND** parameter to preserve partitions.

Saved partitions might be data partitions from an existing OpenShift Container Platform system. You can identify the disk partitions you want to keep either by partition label or by number.



NOTE

If you save existing partitions, and those partitions do not leave enough space for RHCOS, the installation will fail without damaging the saved partitions.

Retaining existing partitions during an ISO installation

This example preserves any partition in which the partition label begins with **data** (**data***):

```
# coreos-installer install --ignition-url http://10.0.2.2:8080/user.ign \
  --save-partlabel 'data*' /dev/disk/by-id/scsi-<serial_number>
```

The following example illustrates running the **coreos-installer** in a way that preserves the sixth (6) partition on the disk:

```
# coreos-installer install --ignition-url http://10.0.2.2:8080/user.ign \
  --save-partindex 6 /dev/disk/by-id/scsi-<serial_number>
```

This example preserves partitions 5 and higher:

```
# coreos-installer install --ignition-url http://10.0.2.2:8080/user.ign \
  --save-partindex 5- /dev/disk/by-id/scsi-<serial_number>
```

In the previous examples where partition saving is used, **coreos-installer** recreates the partition immediately.

Retaining existing partitions during a PXE installation

This **APPEND** option preserves any partition in which the partition label begins with 'data' ('data*'):

```
coreos.inst.save_partlabel=data*
```

This **APPEND** option preserves partitions 5 and higher:

```
coreos.inst.save_partindex=5-
```


This **APPEND** option preserves partition 6:

```
coreos.inst.save_partindex=6
```

15.3.14.3.3. Identifying Ignition configs

When doing an RHCOS manual installation, there are two types of Ignition configs that you can provide, with different reasons for providing each one:

- **Permanent install Ignition config:** Every manual RHCOS installation needs to pass one of the Ignition config files generated by **openshift-installer**, such as **bootstrap.ign**, **master.ign** and **worker.ign**, to carry out the installation.



IMPORTANT

It is not recommended to modify these Ignition config files directly. You can update the manifest files that are wrapped into the Ignition config files, as outlined in examples in the preceding sections.

For PXE installations, you pass the Ignition configs on the **APPEND** line using the **coreos.inst.ignition_url=** option. For ISO installations, after the ISO boots to the shell prompt, you identify the Ignition config on the **coreos-installer** command line with the **--ignition-url=** option. In both cases, only HTTP and HTTPS protocols are supported.

- **Live install Ignition config:** This type can be created by using the **coreos-installer customize** subcommand and its various options. With this method, the Ignition config passes to the live install medium, runs immediately upon booting, and performs setup tasks before or after the RHCOS system installs to disk. This method should only be used for performing tasks that must be done once and not applied again later, such as with advanced partitioning that cannot be done using a machine config.

For PXE or ISO boots, you can create the Ignition config and **APPEND** the **ignition.config.url=** option to identify the location of the Ignition config. You also need to append **ignition.firstboot** **ignition.platform.id=metal** or the **ignition.config.url** option will be ignored.

15.3.14.3.4. Default console configuration

Red Hat Enterprise Linux CoreOS (RHCOS) nodes installed from an OpenShift Container Platform 4.16 boot image use a default console that is meant to accommodate most virtualized and bare metal setups. Different cloud and virtualization platforms may use different default settings depending on the chosen architecture. Bare metal installations use the kernel default settings which typically means the graphical console is the primary console and the serial console is disabled.

The default consoles may not match your specific hardware configuration or you might have specific needs that require you to adjust the default console. For example:

- You want to access the emergency shell on the console for debugging purposes.

- Your cloud platform does not provide interactive access to the graphical console, but provides a serial console.
- You want to enable multiple consoles.

Console configuration is inherited from the boot image. This means that new nodes in existing clusters are unaffected by changes to the default console.

You can configure the console for bare metal installations in the following ways:

- Using **coreos-installer** manually on the command line.
- Using the **coreos-installer iso customize** or **coreos-installer pxe customize** subcommands with the **--dest-console** option to create a custom image that automates the process.



NOTE

For advanced customization, perform console configuration using the **coreos-installer iso** or **coreos-installer pxe** subcommands, and not kernel arguments.

15.3.14.3.5. Enabling the serial console for PXE and ISO installations

By default, the Red Hat Enterprise Linux CoreOS (RHCOS) serial console is disabled and all output is written to the graphical console. You can enable the serial console for an ISO installation and reconfigure the bootloader so that output is sent to both the serial console and the graphical console.

Procedure

1. Boot the ISO installer.
2. Run the **coreos-installer** command to install the system, adding the **--console** option once to specify the graphical console, and a second time to specify the serial console:

```
$ coreos-installer install \
  --console=tty0 \ 1
  --console=ttyS0,<options> \ 2
  --ignition-url=http://host/worker.ign /dev/disk/by-id/scsi-<serial_number>
```

- 1 The desired secondary console. In this case, the graphical console. Omitting this option will disable the graphical console.
- 2 The desired primary console. In this case the serial console. The **options** field defines the baud rate and other settings. A common value for this field is **11520n8**. If no options are provided, the default kernel value of **9600n8** is used. For more information on the format of this option, see [Linux kernel serial console](#) documentation.

3. Reboot into the installed system.



NOTE

A similar outcome can be obtained by using the **coreos-installer install --append-karg** option, and specifying the console with **console=**. However, this will only set the console for the kernel and not the bootloader.

To configure a PXE installation, make sure the **coreos.inst.install_dev** kernel command line option is omitted, and use the shell prompt to run **coreos-installer** manually using the above ISO installation procedure.

15.3.14.3.6. Customizing a live RHCOS ISO or PXE install

You can use the live ISO image or PXE environment to install RHCOS by injecting an Ignition config file directly into the image. This creates a customized image that you can use to provision your system.

For an ISO image, the mechanism to do this is the **coreos-installer iso customize** subcommand, which modifies the **.iso** file with your configuration. Similarly, the mechanism for a PXE environment is the **coreos-installer pxe customize** subcommand, which creates a new **initramfs** file that includes your customizations.

The **customize** subcommand is a general purpose tool that can embed other types of customizations as well. The following tasks are examples of some of the more common customizations:

- Inject custom CA certificates for when corporate security policy requires their use.
- Configure network settings without the need for kernel arguments.
- Embed arbitrary preinstall and post-install scripts or binaries.

15.3.14.3.7. Customizing a live RHCOS ISO image

You can customize a live RHCOS ISO image directly with the **coreos-installer iso customize** subcommand. When you boot the ISO image, the customizations are applied automatically.

You can use this feature to configure the ISO image to automatically install RHCOS.

Procedure

1. Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.
2. Retrieve the RHCOS ISO image from the [RHCOS image mirror](#) page and the Ignition config file, and then run the following command to inject the Ignition config directly into the ISO image:

```
$ coreos-installer iso customize rhcos-<version>-live.x86_64.iso \
  --dest-ignition bootstrap.ign \ 1
  --dest-device /dev/disk/by-id/scsi-<serial_number> 2
```

- 1 The Ignition config file that is generated from the **openshift-installer** installation program.
- 2 When you specify this option, the ISO image automatically runs an installation. Otherwise, the image remains configured for installation, but does not install automatically unless you specify the **coreos.inst.install_dev** kernel argument.

3. Optional: To remove the ISO image customizations and return the image to its pristine state, run:

```
$ coreos-installer iso reset rhcos-<version>-live.x86_64.iso
```

You can now re-customize the live ISO image or use it in its pristine state.

Applying your customizations affects every subsequent boot of RHCOS.

15.3.14.3.7.1. Modifying a live install ISO image to enable the serial console

On clusters installed with OpenShift Container Platform 4.12 and above, the serial console is disabled by default and all output is written to the graphical console. You can enable the serial console with the following procedure.

Procedure

1. Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.
2. Retrieve the RHCOS ISO image from the [RHCOS image mirror](#) page and run the following command to customize the ISO image to enable the serial console to receive output:

```
$ coreos-installer iso customize rhcos-<version>-live.x86_64.iso \
  --dest-ignition <path> \ 1
  --dest-console tty0 \ 2
  --dest-console ttyS0,<options> \ 3
  --dest-device /dev/disk/by-id/scsi-<serial_number> \ 4
```

- 1** The location of the Ignition config to install.
- 2** The desired secondary console. In this case, the graphical console. Omitting this option will disable the graphical console.
- 3** The desired primary console. In this case, the serial console. The **options** field defines the baud rate and other settings. A common value for this field is **115200n8**. If no options are provided, the default kernel value of **9600n8** is used. For more information on the format of this option, see the [Linux kernel serial console](#) documentation.
- 4** The specified disk to install to. If you omit this option, the ISO image automatically runs the installation program which will fail unless you also specify the **coreos.inst.install_dev** kernel argument.



NOTE

The **--dest-console** option affects the installed system and not the live ISO system. To modify the console for a live ISO system, use the **--live-karg-append** option and specify the console with **console=**.

Your customizations are applied and affect every subsequent boot of the ISO image.

3. Optional: To remove the ISO image customizations and return the image to its original state, run the following command:

```
$ coreos-installer iso reset rhcos-<version>-live.x86_64.iso
```

You can now recustomize the live ISO image or use it in its original state.

15.3.14.3.7.2. Modifying a live install ISO image to use a custom certificate authority

You can provide certificate authority (CA) certificates to Ignition with the **--ignition-ca** flag of the **customize** subcommand. You can use the CA certificates during both the installation boot and when provisioning the installed system.



NOTE

Custom CA certificates affect how Ignition fetches remote resources but they do not affect the certificates installed onto the system.

Procedure

1. Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.
2. Retrieve the RHCOS ISO image from the [RHCOS image mirror](#) page and run the following command to customize the ISO image for use with a custom CA:

```
$ coreos-installer iso customize rhcos-<version>-live.x86_64.iso --ignition-ca cert.pem
```



IMPORTANT

The **coreos.inst.ignition_url** kernel parameter does not work with the **--ignition-ca** flag. You must use the **--dest-ignition** flag to create a customized image for each cluster.

Applying your custom CA certificate affects every subsequent boot of RHCOS.

15.3.14.3.7.3. Modifying a live install ISO image with customized network settings

You can embed a NetworkManager keyfile into the live ISO image and pass it through to the installed system with the **--network-keyfile** flag of the **customize** subcommand.



WARNING

When creating a connection profile, you must use a **.nmconnection** filename extension in the filename of the connection profile. If you do not use a **.nmconnection** filename extension, the cluster will apply the connection profile to the live environment, but it will not apply the configuration when the cluster first boots up the nodes, resulting in a setup that does not work.

Procedure

1. Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.
2. Create a connection profile for a bonded interface. For example, create the **bond0.nmconnection** file in your local directory with the following content:

```
[connection]
id=bond0
type=bond
interface-name=bond0
```

```
multi-connect=1
permissions=

[ethernet]
mac-address-blacklist=

[bond]
miimon=100
mode=active-backup

[ipv4]
method=auto

[ipv6]
method=auto

[proxy]
```

3. Create a connection profile for a secondary interface to add to the bond. For example, create the **bond0-proxy-em1.nmconnection** file in your local directory with the following content:

```
[connection]
id=em1
type=ethernet
interface-name=em1
master=bond0
multi-connect=1
permissions=
slave-type=bond

[ethernet]
mac-address-blacklist=
```

4. Create a connection profile for a secondary interface to add to the bond. For example, create the **bond0-proxy-em2.nmconnection** file in your local directory with the following content:

```
[connection]
id=em2
type=ethernet
interface-name=em2
master=bond0
multi-connect=1
permissions=
slave-type=bond

[ethernet]
mac-address-blacklist=
```

5. Retrieve the RHCOS ISO image from the [RHCOS image mirror](#) page and run the following command to customize the ISO image with your configured networking:

```
$ coreos-installer iso customize rhcos-<version>-live.x86_64.iso \
  --network-keyfile bond0.nmconnection \
  --network-keyfile bond0-proxy-em1.nmconnection \
  --network-keyfile bond0-proxy-em2.nmconnection
```


Network settings are applied to the live system and are carried over to the destination system.

15.3.14.3.7.4. Customizing a live install ISO image for an iSCSI boot device

You can set the iSCSI target and initiator values for automatic mounting, booting and configuration using a customized version of the live RHCOS image.

Prerequisites

1. You have an iSCSI target you want to install RHCOS on.

Procedure

1. Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.
2. Retrieve the RHCOS ISO image from the [RHCOS image mirror](#) page and run the following command to customize the ISO image with the following information:

```
$ coreos-installer iso customize \
  --pre-install mount-iscsi.sh \ 1
  --post-install unmount-iscsi.sh \ 2
  --dest-device /dev/disk/by-path/<IP_address>:<port>-iscsi-<target_iqn>-lun-<lun> \ 3
  --dest-ignition config.ign \ 4
  --dest-karg-append rd.iscsi.initiator=<initiator_iqn> \ 5
  --dest-karg-append netroot=<target_iqn> \ 6
  -o custom.iso rhcos-<version>-live.x86_64.iso
```

- 1 The script that gets run before installation. It should contain the **iscsiadm** commands for mounting the iSCSI target and any commands enabling multipathing.
- 2 The script that gets run after installation. It should contain the command **iscsiadm --mode node --logout=all**.
- 3 The location of the destination system. You must provide the IP address of the target portal, the associated port number, the target iSCSI node in IQN format, and the iSCSI logical unit number (LUN).
- 4 The Ignition configuration for the destination system.
- 5 The iSCSI initiator, or client, name in IQN format. The initiator forms a session to connect to the iSCSI target.
- 6 The the iSCSI target, or server, name in IQN format.

For more information about the iSCSI options supported by **dracut**, see the [dracut.cmdline manual page](#).

15.3.14.3.7.5. Customizing a live install ISO image for an iSCSI boot device with iBFT

You can set the iSCSI target and initiator values for automatic mounting, booting and configuration using a customized version of the live RHCOS image.

Prerequisites

1. You have an iSCSI target you want to install RHCOS on.
2. Optional: you have multipathed your iSCSI target.

Procedure

1. Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.
2. Retrieve the RHCOS ISO image from the [RHCOS image mirror](#) page and run the following command to customize the ISO image with the following information:

```
$ coreos-installer iso customize \  
  --pre-install mount-iscsi.sh \ 1  
  --post-install unmount-iscsi.sh \ 2  
  --dest-device /dev/mapper/mpatha \ 3  
  --dest-ignition config.ign \ 4  
  --dest-karg-append rd.iscsi.firmware=1 \ 5  
  --dest-karg-append rd.multipath=default \ 6  
  -o custom.iso rhcos-<version>-live.x86_64.iso
```

- 1 The script that gets run before installation. It should contain the **iscsiadm** commands for mounting the iSCSI target and any commands enabling multipathing.
- 2 The script that gets run after installation. It should contain the command **iscsiadm --mode node --logout=all**.
- 3 The path to the device. If you are using multipath, the multipath device, **/dev/mapper/mpatha**, If there are multiple multipath devices connected, or to be explicit, you can use the World Wide Name (WWN) symlink available in **/dev/disk/by-path**.
- 4 The Ignition configuration for the destination system.
- 5 The iSCSI parameter is read from the BIOS firmware.
- 6 Optional: include this parameter if you are enabling multipathing.

For more information about the iSCSI options supported by **dracut**, see the [dracut.cmdline manual page](#).

15.3.14.3.8. Customizing a live RHCOS PXE environment

You can customize a live RHCOS PXE environment directly with the **coreos-installer pxe customize** subcommand. When you boot the PXE environment, the customizations are applied automatically.

You can use this feature to configure the PXE environment to automatically install RHCOS.

Procedure

1. Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.
2. Retrieve the RHCOS **kernel**, **initramfs** and **rootfs** files from the [RHCOS image mirror](#) page and the Ignition config file, and then run the following command to create a new **initramfs** file that contains the customizations from your Ignition config:

```
$ coreos-installer pxe customize rhcos-<version>-live-initramfs.x86_64.img \
  --dest-ignition bootstrap.ign \ 1
  --dest-device /dev/disk/by-id/scsi-<serial_number> \ 2
  -o rhcos-<version>-custom-initramfs.x86_64.img 3
```

- 1 The Ignition config file that is generated from **openshift-installer**.
- 2 When you specify this option, the PXE environment automatically runs an install. Otherwise, the image remains configured for installing, but does not do so automatically unless you specify the **coreos.inst.install_dev** kernel argument.
- 3 Use the customized **initramfs** file in your PXE configuration. Add the **ignition.firstboot** and **ignition.platform.id=metal** kernel arguments if they are not already present.

Applying your customizations affects every subsequent boot of RHCOS.

15.3.14.3.8.1. Modifying a live install PXE environment to enable the serial console

On clusters installed with OpenShift Container Platform 4.12 and above, the serial console is disabled by default and all output is written to the graphical console. You can enable the serial console with the following procedure.

Procedure

1. Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.
2. Retrieve the RHCOS **kernel**, **initramfs** and **rootfs** files from the [RHCOS image mirror](#) page and the Ignition config file, and then run the following command to create a new customized **initramfs** file that enables the serial console to receive output:

```
$ coreos-installer pxe customize rhcos-<version>-live-initramfs.x86_64.img \
  --dest-ignition <path> \ 1
  --dest-console tty0 \ 2
  --dest-console ttyS0,<options> \ 3
  --dest-device /dev/disk/by-id/scsi-<serial_number> \ 4
  -o rhcos-<version>-custom-initramfs.x86_64.img 5
```

- 1 The location of the Ignition config to install.
- 2 The desired secondary console. In this case, the graphical console. Omitting this option will disable the graphical console.
- 3 The desired primary console. In this case, the serial console. The **options** field defines the baud rate and other settings. A common value for this field is **115200n8**. If no options are provided, the default kernel value of **9600n8** is used. For more information on the format of this option, see the [Linux kernel serial console](#) documentation.
- 4 The specified disk to install to. If you omit this option, the PXE environment automatically runs the installer which will fail unless you also specify the **coreos.inst.install_dev** kernel argument.
- 5 Use the customized **initramfs** file in your PXE configuration. Add the **ignition.firstboot** and **ignition.platform.id=metal** kernel arguments if they are not already present.

Your customizations are applied and affect every subsequent boot of the PXE environment.

15.3.14.3.8.2. Modifying a live install PXE environment to use a custom certificate authority

You can provide certificate authority (CA) certificates to Ignition with the **--ignition-ca** flag of the **customize** subcommand. You can use the CA certificates during both the installation boot and when provisioning the installed system.



NOTE


Custom CA certificates affect how Ignition fetches remote resources but they do not affect the certificates installed onto the system.

Procedure

1. Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.
2. Retrieve the RHCOS **kernel**, **initramfs** and **rootfs** files from the [RHCOS image mirror](#) page and run the following command to create a new customized **initramfs** file for use with a custom CA:

```
$ coreos-installer pxe customize rhcos-<version>-live-initramfs.x86_64.img \
--ignition-ca cert.pem \
-o rhcos-<version>-custom-initramfs.x86_64.img
```

3. Use the customized **initramfs** file in your PXE configuration. Add the **ignition.firstboot** and **ignition.platform.id=metal** kernel arguments if they are not already present.



IMPORTANT

The **coreos.inst.ignition_url** kernel parameter does not work with the **--ignition-ca** flag. You must use the **--dest-ignition** flag to create a customized image for each cluster.

Applying your custom CA certificate affects every subsequent boot of RHCOS.

15.3.14.3.8.3. Modifying a live install PXE environment with customized network settings

You can embed a NetworkManager keyfile into the live PXE environment and pass it through to the installed system with the **--network-keyfile** flag of the **customize** subcommand.



WARNING

When creating a connection profile, you must use a **.nmconnection** filename extension in the filename of the connection profile. If you do not use a **.nmconnection** filename extension, the cluster will apply the connection profile to the live environment, but it will not apply the configuration when the cluster first boots up the nodes, resulting in a setup that does not work.

Procedure

1. Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.
2. Create a connection profile for a bonded interface. For example, create the **bond0.nmconnection** file in your local directory with the following content:

```
[connection]
id=bond0
type=bond
interface-name=bond0
multi-connect=1
permissions=

[ethernet]
mac-address-blacklist=

[bond]
miimon=100
mode=active-backup

[ipv4]
method=auto

[ipv6]
method=auto

[proxy]
```

3. Create a connection profile for a secondary interface to add to the bond. For example, create the **bond0-proxy-em1.nmconnection** file in your local directory with the following content:

```
[connection]
id=em1
type=ethernet
interface-name=em1
master=bond0
multi-connect=1
permissions=
slave-type=bond

[ethernet]
mac-address-blacklist=
```

4. Create a connection profile for a secondary interface to add to the bond. For example, create the **bond0-proxy-em2.nmconnection** file in your local directory with the following content:

```
[connection]
id=em2
type=ethernet
interface-name=em2
master=bond0
multi-connect=1
permissions=
slave-type=bond
```

```
[ethernet]
mac-address-blacklist=
```

- Retrieve the RHCOS **kernel**, **initramfs** and **rootfs** files from the [RHCOS image mirror](#) page and run the following command to create a new customized **initramfs** file that contains your configured networking:

```
$ coreos-installer pxe customize rhcos-<version>-live-initramfs.x86_64.img \
  --network-keyfile bond0.nmconnection \
  --network-keyfile bond0-proxy-em1.nmconnection \
  --network-keyfile bond0-proxy-em2.nmconnection \
  -o rhcos-<version>-custom-initramfs.x86_64.img
```

- Use the customized **initramfs** file in your PXE configuration. Add the **ignition.firstboot** and **ignition.platform.id=metal** kernel arguments if they are not already present. Network settings are applied to the live system and are carried over to the destination system.

15.3.14.3.8.4. Customizing a live install PXE environment for an iSCSI boot device

You can set the iSCSI target and initiator values for automatic mounting, booting and configuration using a customized version of the live RHCOS image.

Prerequisites

- You have an iSCSI target you want to install RHCOS on.

Procedure

- Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.
- Retrieve the RHCOS **kernel**, **initramfs** and **rootfs** files from the [RHCOS image mirror](#) page and run the following command to create a new customized **initramfs** file with the following information:

```
$ coreos-installer pxe customize \
  --pre-install mount-iscsi.sh \ 1
  --post-install unmount-iscsi.sh \ 2
  --dest-device /dev/disk/by-path/<IP_address>:<port>-iscsi-<target_iqn>-lun-<lun> \ 3
  --dest-ignition config.ign \ 4
  --dest-karg-append rd.iscsi.initiator=<initiator_iqn> \ 5
  --dest-karg-append netroot=<target_iqn> \ 6
  -o custom.img rhcos-<version>-live-initramfs.x86_64.img
```

- The script that gets run before installation. It should contain the **iscsiadm** commands for mounting the iSCSI target and any commands enabling multipathing.
- The script that gets run after installation. It should contain the command **iscsiadm --mode node --logout=all**.
- The location of the destination system. You must provide the IP address of the target portal, the associated port number, the target iSCSI node in IQN format, and the iSCSI logical unit number (LUN).

- 4 The Ignition configuration for the destination system.
- 5 The iSCSI initiator, or client, name in IQN format. The initiator forms a session to connect to the iSCSI target.
- 6 The the iSCSI target, or server, name in IQN format.

For more information about the iSCSI options supported by **dracut**, see the [dracut.cmdline manual page](#).

15.3.14.3.8.5. Customizing a live install PXE environment for an iSCSI boot device with iBFT

You can set the iSCSI target and initiator values for automatic mounting, booting and configuration using a customized version of the live RHCOS image.

Prerequisites

1. You have an iSCSI target you want to install RHCOS on.
2. Optional: you have multipathed your iSCSI target.

Procedure

1. Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.
2. Retrieve the RHCOS **kernel**, **initramfs** and **rootfs** files from the [RHCOS image mirror](#) page and run the following command to create a new customized **initramfs** file with the following information:

```
$ coreos-installer pxe customize \
  --pre-install mount-iscsi.sh \ 1
  --post-install unmount-iscsi.sh \ 2
  --dest-device /dev/mapper/mpatha \ 3
  --dest-ignition config.ign \ 4
  --dest-karg-append rd.iscsi.firmware=1 \ 5
  --dest-karg-append rd.multipath=default \ 6
  -o custom.img rhcos-<version>-live-initramfs.x86_64.img
```

- 1 The script that gets run before installation. It should contain the **iscsiadm** commands for mounting the iSCSI target.
- 2 The script that gets run after installation. It should contain the command **iscsiadm --mode node --logout=all**.
- 3 The path to the device. If you are using multipath, the multipath device, **/dev/mapper/mpatha**, If there are multiple multipath devices connected, or to be explicit, you can use the World Wide Name (WWN) symlink available in **/dev/disk/by-path**.
- 4 The Ignition configuration for the destination system.
- 5 The iSCSI parameter is read from the BIOS firmware.
- 6 Optional: include this parameter if you are enabling multipathing.


For more information about the iSCSI options supported by **dracut**, see the [dracut.cmdline manual page](#).

15.3.14.3.9. Advanced RHCOS installation reference

This section illustrates the networking configuration and other advanced options that allow you to modify the Red Hat Enterprise Linux CoreOS (RHCOS) manual installation process. The following tables describe the kernel arguments and command-line options you can use with the RHCOS live installer and the **coreos-installer** command.

15.3.14.3.9.1. Networking and bonding options for ISO installations

If you install RHCOS from an ISO image, you can add kernel arguments manually when you boot the image to configure networking for a node. If no networking arguments are specified, DHCP is activated in the initramfs when RHCOS detects that networking is required to fetch the Ignition config file.



IMPORTANT

When adding networking arguments manually, you must also add the **rd.neednet=1** kernel argument to bring the network up in the initramfs.

The following information provides examples for configuring networking and bonding on your RHCOS nodes for ISO installations. The examples describe how to use the **ip=**, **nameserver=**, and **bond=** kernel arguments.



NOTE

Ordering is important when adding the kernel arguments: **ip=**, **nameserver=**, and then **bond=**.

The networking options are passed to the **dracut** tool during system boot. For more information about the networking options supported by **dracut**, see the [dracut.cmdline manual page](#).

The following examples are the networking options for ISO installation.

Configuring DHCP or static IP addresses

To configure an IP address, either use DHCP (**ip=dhcp**) or set an individual static IP address (**ip=<host_ip>**). If setting a static IP, you must then identify the DNS server IP address (**nameserver=<dns_ip>**) on each node. The following example sets:

- The node's IP address to **10.10.10.2**
- The gateway address to **10.10.10.254**
- The netmask to **255.255.255.0**
- The hostname to **core0.example.com**
- The DNS server address to **4.4.4.41**
- The auto-configuration value to **none**. No auto-configuration is required when IP networking is configured statically.


```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none
nameserver=4.4.4.41
```



NOTE

When you use DHCP to configure IP addressing for the RHCOS machines, the machines also obtain the DNS server information through DHCP. For DHCP-based deployments, you can define the DNS server address that is used by the RHCOS nodes through your DHCP server configuration.

Configuring an IP address without a static hostname

You can configure an IP address without assigning a static hostname. If a static hostname is not set by the user, it will be picked up and automatically set by a reverse DNS lookup. To configure an IP address without a static hostname refer to the following example:

- The node's IP address to **10.10.10.2**
- The gateway address to **10.10.10.254**
- The netmask to **255.255.255.0**
- The DNS server address to **4.4.4.41**
- The auto-configuration value to **none**. No auto-configuration is required when IP networking is configured statically.

```
ip=10.10.10.2::10.10.10.254:255.255.255.0::enp1s0:none
nameserver=4.4.4.41
```

Specifying multiple network interfaces

You can specify multiple network interfaces by setting multiple **ip=** entries.

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none
ip=10.10.10.3::10.10.10.254:255.255.255.0:core0.example.com:enp2s0:none
```

Configuring default gateway and route

Optional: You can configure routes to additional networks by setting an **rd.route=** value.



NOTE

When you configure one or multiple networks, one default gateway is required. If the additional network gateway is different from the primary network gateway, the default gateway must be the primary network gateway.

- Run the following command to configure the default gateway:

```
ip=::10.10.10.254:::
```

- Enter the following command to configure the route for the additional network:

```
rd.route=20.20.20.0/24:20.20.20.254:enp2s0
```

Disabling DHCP on a single interface

You can disable DHCP on a single interface, such as when there are two or more network interfaces and only one interface is being used. In the example, the **enp1s0** interface has a static networking configuration and DHCP is disabled for **enp2s0**, which is not used:

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none
ip=::::core0.example.com:enp2s0:none
```

Combining DHCP and static IP configurations

You can combine DHCP and static IP configurations on systems with multiple network interfaces, for example:

```
ip=enp1s0:dhcp
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp2s0:none
```

Configuring VLANs on individual interfaces

Optional: You can configure VLANs on individual interfaces by using the **vlan=** parameter.

- To configure a VLAN on a network interface and use a static IP address, run the following command:

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp2s0.100:none
vlan=enp2s0.100:enp2s0
```

- To configure a VLAN on a network interface and to use DHCP, run the following command:

```
ip=enp2s0.100:dhcp
vlan=enp2s0.100:enp2s0
```

Providing multiple DNS servers

You can provide multiple DNS servers by adding a **nameserver=** entry for each server, for example:

```
nameserver=1.1.1.1
nameserver=8.8.8.8
```

Bonding multiple network interfaces to a single interface

Optional: You can bond multiple network interfaces to a single interface by using the **bond=** option. Refer to the following examples:


- The syntax for configuring a bonded interface is: **bond=<name>[:<network_interfaces>][:options]**
<name> is the bonding device name (**bond0**), **<network_interfaces>** represents a comma-separated list of physical (ethernet) interfaces (**em1,em2**), and *options* is a comma-separated list of bonding options. Enter **modinfo bonding** to see available options.
- When you create a bonded interface using **bond=**, you must specify how the IP address is assigned and other information for the bonded interface.
 - To configure the bonded interface to use DHCP, set the bond's IP address to **dhcp**. For example:

```
bond=bond0:em1,em2:mode=active-backup
ip=bond0:dhcp
```

- To configure the bonded interface to use a static IP address, enter the specific IP address you want and related information. For example:

```
bond=bond0:em1,em2:mode=active-backup
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:bond0:none
```

Bonding multiple SR-IOV network interfaces to a dual port NIC interface



IMPORTANT

Support for Day 1 operations associated with enabling NIC partitioning for SR-IOV devices is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

Optional: You can bond multiple SR-IOV network interfaces to a dual port NIC interface by using the **bond=** option.

On each node, you must perform the following tasks:

1. Create the SR-IOV virtual functions (VFs) following the guidance in [Managing SR-IOV devices](#). Follow the procedure in the "Attaching SR-IOV networking devices to virtual machines" section.
2. Create the bond, attach the desired VFs to the bond and set the bond link state up following the guidance in [Configuring network bonding](#). Follow any of the described procedures to create the bond.

The following examples illustrate the syntax you must use:

- The syntax for configuring a bonded interface is **bond=<name>[:<network_interfaces>][:options]**. **<name>** is the bonding device name (**bond0**), **<network_interfaces>** represents the virtual functions (VFs) by their known name in the kernel and shown in the output of the **ip link** command (**eno1f0**, **eno2f0**), and **options** is a comma-separated list of bonding options. Enter **modinfo bonding** to see available options.
- When you create a bonded interface using **bond=**, you must specify how the IP address is assigned and other information for the bonded interface.
 - To configure the bonded interface to use DHCP, set the bond's IP address to **dhcp**. For example:

```
bond=bond0:eno1f0,eno2f0:mode=active-backup
ip=bond0:dhcp
```

- To configure the bonded interface to use a static IP address, enter the specific IP address you want and related information. For example:

```
bond=bond0:eno1f0,eno2f0:mode=active-backup
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:bond0:none
```

Using network teaming

Optional: You can use a network teaming as an alternative to bonding by using the **team=** parameter:

- The syntax for configuring a team interface is: **team=name[:network_interfaces]**
name is the team device name (**team0**) and *network_interfaces* represents a comma-separated list of physical (ethernet) interfaces (**em1, em2**).



NOTE

Teaming is planned to be deprecated when RHCOS switches to an upcoming version of RHEL. For more information, see this [Red Hat Knowledgebase Article](#).

Use the following example to configure a network team:

```
team=team0:em1,em2
ip=team0:dhcp
```


15.3.14.3.9.2. coreos-installer options for ISO and PXE installations

You can install RHCOS by running **coreos-installer install <options> <device>** at the command prompt, after booting into the RHCOS live environment from an ISO image.

The following table shows the subcommands, options, and arguments you can pass to the **coreos-installer** command.


Table 15.30. coreos-installer subcommands, command-line options, and arguments

coreos-installer install subcommand	
<i>Subcommand</i>	<i>Description</i>
\$ coreos-installer install <options> <device>	Embed an Ignition config in an ISO image.
coreos-installer install subcommand options	
<i>Option</i>	<i>Description</i>
-u, --image-url <url>	Specify the image URL manually.
-f, --image-file <path>	Specify a local image file manually. Used for debugging.
-i, --ignition-file <path>	Embed an Ignition config from a file.
-l, --ignition-url <URL>	Embed an Ignition config from a URL.
--ignition-hash <digest>	Digest type-value of the Ignition config.
-p, --platform <name>	Override the Ignition platform ID for the installed system.

--console <spec>	Set the kernel and bootloader console for the installed system. For more information about the format of <spec> , see the Linux kernel serial console documentation.
--append-karg <arg>...	Append a default kernel argument to the installed system.
--delete-karg <arg>...	Delete a default kernel argument from the installed system.
-n, --copy-network	Copy the network configuration from the install environment. <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p>IMPORTANT</p> <p>The --copy-network option only copies networking configuration found under /etc/NetworkManager/system-connections. In particular, it does not copy the system hostname.</p> </div> </div>
--network-dir <path>	For use with -n . Default is /etc/NetworkManager/system-connections/ .
--save-partlabel <lx>..	Save partitions with this label glob.
--save-partindex <id>...	Save partitions with this number or range.
--insecure	Skip RHCOS image signature verification.
--insecure-ignition	Allow Ignition URL without HTTPS or hash.
--architecture <name>	Target CPU architecture. Valid values are x86_64 and aarch64 .
--preserve-on-error	Do not clear partition table on error.
-h, --help	Print help information.
coreos-installer install subcommand argument	
<i>Argument</i>	<i>Description</i>
<device>	The destination device.
coreos-installer ISO subcommands	

<i>Subcommand</i>	<i>Description</i>
\$ coreos-installer iso customize <options> <ISO_image>	Customize a RHCOS live ISO image.
coreos-installer iso reset <options> <ISO_image>	Restore a RHCOS live ISO image to default settings.
coreos-installer iso ignition remove <options> <ISO_image>	Remove the embedded Ignition config from an ISO image.
coreos-installer ISO customize subcommand options	
<i>Option</i>	<i>Description</i>
--dest-ignition <path>	Merge the specified Ignition config file into a new configuration fragment for the destination system.
--dest-console <spec>	Specify the kernel and bootloader console for the destination system.
--dest-device <path>	Install and overwrite the specified destination device.
--dest-karg-append <arg>	Add a kernel argument to each boot of the destination system.
--dest-karg-delete <arg>	Delete a kernel argument from each boot of the destination system.
--network-keyfile <path>	Configure networking by using the specified NetworkManager keyfile for live and destination systems.
--ignition-ca <path>	Specify an additional TLS certificate authority to be trusted by Ignition.
--pre-install <path>	Run the specified script before installation.
--post-install <path>	Run the specified script after installation.
--installer-config <path>	Apply the specified installer configuration file.
--live-ignition <path>	Merge the specified Ignition config file into a new configuration fragment for the live environment.
--live-karg-append <arg>	Add a kernel argument to each boot of the live environment.

--live-karg-delete <arg>	Delete a kernel argument from each boot of the live environment.
--live-karg-replace <k=o=n>	Replace a kernel argument in each boot of the live environment, in the form key=old=new .
-f, --force	Overwrite an existing Ignition config.
-o, --output <path>	Write the ISO to a new output file.
-h, --help	Print help information.
coreos-installer PXE subcommands	
<i>Subcommand</i>	<i>Description</i>
Note that not all of these options are accepted by all subcommands.	
coreos-installer pxe customize <options> <path>	Customize a RHCOS live PXE boot config.
coreos-installer pxe ignition wrap <options>	Wrap an Ignition config in an image.
coreos-installer pxe ignition unwrap <options> <image_name>	Show the wrapped Ignition config in an image.
coreos-installer PXE customize subcommand options	
<i>Option</i>	<i>Description</i>
Note that not all of these options are accepted by all subcommands.	
--dest-ignition <path>	Merge the specified Ignition config file into a new configuration fragment for the destination system.
--dest-console <spec>	Specify the kernel and bootloader console for the destination system.
--dest-device <path>	Install and overwrite the specified destination device.
--network-keyfile <path>	Configure networking by using the specified NetworkManager keyfile for live and destination systems.
--ignition-ca <path>	Specify an additional TLS certificate authority to be trusted by Ignition.
--pre-install <path>	Run the specified script before installation.

post-install <path>	Run the specified script after installation.
--installer-config <path>	Apply the specified installer configuration file.
--live-ignition <path>	Merge the specified Ignition config file into a new configuration fragment for the live environment.
-o, --output <path>	Write the initramfs to a new output file.
	 <p>NOTE</p> <p>This option is required for PXE environments.</p>
-h, --help	Print help information.

15.3.14.3.9.3. coreos.inst boot options for ISO or PXE installations

You can automatically invoke **coreos-installer** options at boot time by passing **coreos.inst** boot arguments to the RHCOS live installer. These are provided in addition to the standard boot arguments.

- For ISO installations, the **coreos.inst** options can be added by interrupting the automatic boot at the bootloader menu. You can interrupt the automatic boot by pressing **TAB** while the **RHEL CoreOS (Live)** menu option is highlighted.
- For PXE or iPXE installations, the **coreos.inst** options must be added to the **APPEND** line before the RHCOS live installer is booted.

The following table shows the RHCOS live installer **coreos.inst** boot options for ISO and PXE installations.

Table 15.31. **coreos.inst** boot options

Argument	Description
coreos.inst.install_dev	Required. The block device on the system to install to. It is recommended to use the full path, such as /dev/sda , although sda is allowed.
coreos.inst.ignition_url	Optional: The URL of the Ignition config to embed into the installed system. If no URL is specified, no Ignition config is embedded. Only HTTP and HTTPS protocols are supported.
coreos.inst.save_partlabel	Optional: Comma-separated labels of partitions to preserve during the install. Glob-style wildcards are permitted. The specified partitions do not need to exist.

Argument	Description
coreos.inst.save_partindex	Optional: Comma-separated indexes of partitions to preserve during the install. Ranges m-n are permitted, and either m or n can be omitted. The specified partitions do not need to exist.
coreos.inst.insecure	Optional: Permits the OS image that is specified by coreos.inst.image_url to be unsigned.
coreos.inst.image_url	<p>Optional: Download and install the specified RHCOS image.</p> <ul style="list-style-type: none"> ● This argument should not be used in production environments and is intended for debugging purposes only. ● While this argument can be used to install a version of RHCOS that does not match the live media, it is recommended that you instead use the media that matches the version you want to install. ● If you are using coreos.inst.image_url, you must also use coreos.inst.insecure. This is because the bare-metal media are not GPG-signed for OpenShift Container Platform. ● Only HTTP and HTTPS protocols are supported.
coreos.inst.skip_reboot	Optional: The system will not reboot after installing. After the install finishes, you will receive a prompt that allows you to inspect what is happening during installation. This argument should not be used in production environments and is intended for debugging purposes only.
coreos.inst.platform_id	Optional: The Ignition platform ID of the platform the RHCOS image is being installed on. Default is metal . This option determines whether or not to request an Ignition config from the cloud provider, such as VMware. For example: coreos.inst.platform_id=vmware .
ignition.config.url	Optional: The URL of the Ignition config for the live boot. For example, this can be used to customize how coreos-installer is invoked, or to run code before or after the installation. This is different from coreos.inst.ignition_url , which is the Ignition config for the installed system.

15.3.14.4. Enabling multipathing with kernel arguments on RHCOS

RHCOS supports multipathing on the primary disk, allowing stronger resilience to hardware failure to achieve higher host availability.

You can enable multipathing at installation time for nodes that were provisioned in OpenShift Container Platform 4.8 or later. While postinstallation support is available by activating multipathing via the machine config, enabling multipathing during installation is recommended.

In setups where any I/O to non-optimized paths results in I/O system errors, you must enable multipathing at installation time.



IMPORTANT

On IBM Z® and IBM® LinuxONE, you can enable multipathing only if you configured your cluster for it during installation. For more information, see "Installing RHCOS and starting the OpenShift Container Platform bootstrap process" in *Installing a cluster with z/VM on IBM Z® and IBM® LinuxONE*.

The following procedure enables multipath at installation time and appends kernel arguments to the **coreos-installer install** command so that the installed system itself will use multipath beginning from the first boot.



NOTE

OpenShift Container Platform does not support enabling multipathing as a day-2 activity on nodes that have been upgraded from 4.6 or earlier.

Procedure

1. To enable multipath and start the **multipathd** daemon, run the following command on the installation host:

```
$ mpathconf --enable && systemctl start multipathd.service
```

- Optional: If booting the PXE or ISO, you can instead enable multipath by adding **rd.multipath=default** from the kernel command line.

2. Append the kernel arguments by invoking the **coreos-installer** program:

- If there is only one multipath device connected to the machine, it should be available at path **/dev/mapper/mpatha**. For example:

```
$ coreos-installer install /dev/mapper/mpatha \ 1
--append-karg rd.multipath=default \
--append-karg root=/dev/disk/by-label/dm-mpath-root \
--append-karg rw
```

- 1 Indicates the path of the single multipathed device.

- If there are multiple multipath devices connected to the machine, or to be more explicit, instead of using **/dev/mapper/mpatha**, it is recommended to use the World Wide Name (WWN) symlink available in **/dev/disk/by-id**. For example:

```
$ coreos-installer install /dev/disk/by-id/wwn-<wwn_ID> \ 1
--append-karg rd.multipath=default \
--append-karg root=/dev/disk/by-label/dm-mpath-root \
--append-karg rw
```

- 1** Indicates the WWN ID of the target multipathed device. For example, **0xx194e957fcedb4841**.

This symlink can also be used as the **coreos.inst.install_dev** kernel argument when using special **coreos.inst.*** arguments to direct the live installer. For more information, see "Installing RHCOS and starting the OpenShift Container Platform bootstrap process".

3. Check that the kernel arguments worked by going to one of the worker nodes and listing the kernel command line arguments (in **/proc/cmdline** on the host):

```
$ oc debug node/ip-10-0-141-105.ec2.internal
```

Example output

```
Starting pod/ip-10-0-141-105ec2internal-debug ...
To use host binaries, run `chroot /host`

sh-4.2# cat /host/proc/cmdline
...
rd.multipath=default root=/dev/disk/by-label/dm-mpath-root
...

sh-4.2# exit
```

You should see the added kernel arguments.

15.3.14.5. Installing RHCOS manually on an iSCSI boot device

You can manually install RHCOS on an iSCSI target.

Prerequisites

1. You are in the RHCOS live environment.
2. You have an iSCSI target that you want to install RHCOS on.

Procedure

1. Mount the iSCSI target from the live environment by running the following command:

```
$ iscsiadm \
--mode discovery \
--type sendtargets
--portal <IP_address> \ 1
--login
```

- 1** The IP address of the target portal.

2. Install RHCOS onto the iSCSI target by running the following command and using the necessary kernel arguments, for example:

```
$ coreos-installer install \
  /dev/disk/by-path/ip-<IP_address>:<port>-iscsi-<target_iqn>-lun-<lun> \ 1
  --append-karg rd.iscsi.initiator=<initiator_iqn> \ 2
  --append.karg netroot=<target_iqn> \ 3
  --console ttyS0,115200n8
  --ignition-file <path_to_file>
```

- 1** The location you are installing to. You must provide the IP address of the target portal, the associated port number, the target iSCSI node in IQN format, and the iSCSI logical unit number (LUN).
- 2** The iSCSI initiator, or client, name in IQN format. The initiator forms a session to connect to the iSCSI target.
- 3** The the iSCSI target, or server, name in IQN format.

For more information about the iSCSI options supported by **dracut**, see the [dracut.cmdline manual page](#).

3. Unmount the iSCSI disk with the following command:

```
$ iscsiadm --mode node --logoutall=all
```

This procedure can also be performed using the **coreos-installer iso customize** or **coreos-installer pxe customize** subcommands.

15.3.14.6. Installing RHCOS on an iSCSI boot device using iBFT

On a completely diskless machine, the iSCSI target and initiator values can be passed through iBFT. iSCSI multipathing is also supported.

Prerequisites

1. You are in the RHCOS live environment.
2. You have an iSCSI target you want to install RHCOS on.
3. Optional: you have multipathed your iSCSI target.

Procedure

1. Mount the iSCSI target from the live environment by running the following command:

```
$ iscsiadm \
  --mode discovery \
  --type sendtargets \
  --portal <IP_address> \ 1
  --login
```

- 1** The IP address of the target portal.

- Optional: enable multipathing and start the daemon with the following command:

```
$ mpathconf --enable && systemctl start multipathd.service
```

- Install RHCOS onto the iSCSI target by running the following command and using the necessary kernel arguments, for example:

```
$ coreos-installer install \
  /dev/mapper/mpatha \ 1
  --append-karg rd.iscsi.firmware=1 \ 2
  --append-karg rd.multipath=default \ 3
  --console ttyS0 \
  --ignition-file <path_to_file>
```

1 The path of a single multipathed device. If there are multiple multipath devices connected, or to be explicit, you can use the World Wide Name (WWN) symlink available in **/dev/disk/by-path**.

2 The iSCSI parameter is read from the BIOS firmware.

3 Optional: include this parameter if you are enabling multipathing.

For more information about the iSCSI options supported by **dracut**, see the [dracut.cmdline manual page](#).

- Unmount the iSCSI disk:

```
$ iscsiadm --mode node --logout=all
```

This procedure can also be performed using the **coreos-installer iso customize** or **coreos-installer pxe customize** subcommands.

15.3.15. Waiting for the bootstrap process to complete

The OpenShift Container Platform bootstrap process begins after the cluster nodes first boot into the persistent RHCOS environment that has been installed to disk. The configuration information provided through the Ignition config files is used to initialize the bootstrap process and install OpenShift Container Platform on the machines. You must wait for the bootstrap process to complete.

Prerequisites

- You have created the Ignition config files for your cluster.
- You have configured suitable network, DNS and load balancing infrastructure.
- You have obtained the installation program and generated the Ignition config files for your cluster.
- You installed RHCOS on your cluster machines and provided the Ignition config files that the OpenShift Container Platform installation program generated.
- Your machines have direct internet access or have an HTTP or HTTPS proxy available.

Procedure

1. Monitor the bootstrap process:

```
$ ./openshift-install --dir <installation_directory> wait-for bootstrap-complete \ 1
--log-level=info 2
```

- 1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.
- 2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Example output

```
INFO Waiting up to 30m0s for the Kubernetes API at https://api.test.example.com:6443...
INFO API v1.29.4 up
INFO Waiting up to 30m0s for bootstrapping to complete...
INFO It is now safe to remove the bootstrap resources
```

The command succeeds when the Kubernetes API server signals that it has been bootstrapped on the control plane machines.

2. After the bootstrap process is complete, remove the bootstrap machine from the load balancer.



IMPORTANT

You must remove the bootstrap machine from the load balancer at this point. You can also remove or reformat the bootstrap machine itself.

Additional resources

- See [Monitoring installation progress](#) for more information about monitoring the installation logs and retrieving diagnostic data if installation issues arise.

15.3.16. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1 For `<installation_directory>`, specify the path to the directory that you stored the installation files in.

2. Verify you can run `oc` commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

15.3.17. Approving the certificate signing requests for your machines

When you add machines to a cluster, two pending certificate signing requests (CSRs) are generated for each machine that you added. You must confirm that these CSRs are approved or, if necessary, approve them yourself. The client requests must be approved first, followed by the server requests.

Prerequisites

- You added machines to your cluster.

Procedure


1. Confirm that the cluster recognizes the machines:

```
$ oc get nodes
```

Example output

```
NAME      STATUS  ROLES  AGE  VERSION
master-0  Ready   master 63m  v1.29.4
master-1  Ready   master 63m  v1.29.4
master-2  Ready   master 64m  v1.29.4
```

The output lists all of the machines that you created.



NOTE

The preceding output might not include the compute nodes, also known as worker nodes, until some CSRs are approved.

2. Review the pending CSRs and ensure that you see the client requests with the **Pending** or **Approved** status for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

```
NAME      AGE  REQUESTOR                                     CONDITION
csr-8b2br 15m  system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper Pending
```

```
csr-8vnps 15m system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper Pending
...
```

In this example, two machines are joining the cluster. You might see more approved CSRs in the list.

- If the CSRs were not approved, after all of the pending CSRs for the machines you added are in **Pending** status, approve the CSRs for your cluster machines:



NOTE

Because the CSRs rotate automatically, approve your CSRs within an hour of adding the machines to the cluster. If you do not approve them within an hour, the certificates will rotate, and more than two certificates will be present for each node. You must approve all of these certificates. After the client CSR is approved, the Kubelet creates a secondary CSR for the serving certificate, which requires manual approval. Then, subsequent serving certificate renewal requests are automatically approved by the **machine-approver** if the Kubelet requests a new certificate with identical parameters.



NOTE

For clusters running on platforms that are not machine API enabled, such as bare metal and other user-provisioned infrastructure, you must implement a method of automatically approving the kubelet serving certificate requests (CSRs). If a request is not approved, then the **oc exec**, **oc rsh**, and **oc logs** commands cannot succeed, because a serving certificate is required when the API server connects to the kubelet. Any operation that contacts the Kubelet endpoint requires this certificate approval to be in place. The method must watch for new CSRs, confirm that the CSR was submitted by the **node-bootstrapper** service account in the **system:node** or **system:admin** groups, and confirm the identity of the node.

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{"\n"}
{{end}}{{end}}' | xargs --no-run-if-empty oc adm certificate approve
```



NOTE

Some Operators might not become available until some CSRs are approved.

- Now that your client requests are approved, you must review the server requests for each machine that you added to the cluster:


```
$ oc get csr
```

Example output

```
NAME      AGE    REQUESTOR                                CONDITION
csr-bfd72 5m26s system:node:ip-10-0-50-126.us-east-2.compute.internal
Pending
csr-c57lv 5m26s system:node:ip-10-0-95-157.us-east-2.compute.internal
Pending
...
```

- If the remaining CSRs are not approved, and are in the **Pending** status, approve the CSRs for your cluster machines:

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:


```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{"\n"}\n{{end}}{{end}}' | xargs oc adm certificate approve
```

- After all client and server CSRs have been approved, the machines have the **Ready** status. Verify this by running the following command:

```
$ oc get nodes
```

Example output

```
NAME      STATUS  ROLES  AGE  VERSION
master-0  Ready   master 73m  v1.29.4
master-1  Ready   master 73m  v1.29.4
master-2  Ready   master 74m  v1.29.4
worker-0  Ready   worker 11m  v1.29.4
worker-1  Ready   worker 11m  v1.29.4
```



NOTE

It can take a few minutes after approval of the server CSRs for the machines to transition to the **Ready** status.

Additional information

- For more information on CSRs, see [Certificate Signing Requests](#).

15.3.18. Initial Operator configuration

After the control plane initializes, you must immediately configure some Operators so that they all become available.

Prerequisites

- Your control plane has initialized.

Procedure

1. Watch the cluster components come online:

```
$ watch -n5 oc get clusteroperators
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.16.0	True	False	False	19m
baremetal	4.16.0	True	False	False	37m
cloud-credential	4.16.0	True	False	False	40m
cluster-autoscaler	4.16.0	True	False	False	37m
config-operator	4.16.0	True	False	False	38m
console	4.16.0	True	False	False	26m
csi-snapshot-controller	4.16.0	True	False	False	37m
dns	4.16.0	True	False	False	37m
etcd	4.16.0	True	False	False	36m
image-registry	4.16.0	True	False	False	31m
ingress	4.16.0	True	False	False	30m
insights	4.16.0	True	False	False	31m
kube-apiserver	4.16.0	True	False	False	26m
kube-controller-manager	4.16.0	True	False	False	36m
kube-scheduler	4.16.0	True	False	False	36m
kube-storage-version-migrator	4.16.0	True	False	False	37m
machine-api	4.16.0	True	False	False	29m
machine-approver	4.16.0	True	False	False	37m
machine-config	4.16.0	True	False	False	36m
marketplace	4.16.0	True	False	False	37m
monitoring	4.16.0	True	False	False	29m
network	4.16.0	True	False	False	38m
node-tuning	4.16.0	True	False	False	37m
openshift-apiserver	4.16.0	True	False	False	32m
openshift-controller-manager	4.16.0	True	False	False	30m
openshift-samples	4.16.0	True	False	False	32m
operator-lifecycle-manager	4.16.0	True	False	False	37m
operator-lifecycle-manager-catalog	4.16.0	True	False	False	37m
operator-lifecycle-manager-packageserver	4.16.0	True	False	False	32m
service-ca	4.16.0	True	False	False	38m
storage	4.16.0	True	False	False	37m

2. Configure the Operators that are not available.

Additional resources

- See [Gathering logs from a failed installation](#) for details about gathering data in the event of a failed OpenShift Container Platform installation.
- See [Troubleshooting Operator issues](#) for steps to check Operator pod health across the cluster and gather Operator logs for diagnosis.

15.3.18.1. Image registry removed during installation

On platforms that do not provide shareable object storage, the OpenShift Image Registry Operator bootstraps itself as **Removed**. This allows **openshift-installer** to complete installations on these platform types.

After installation, you must edit the Image Registry Operator configuration to switch the **managementState** from **Removed** to **Managed**. When this has completed, you must configure storage.

15.3.18.2. Image registry storage configuration

The Image Registry Operator is not initially available for platforms that do not provide default storage. After installation, you must configure your registry to use storage so that the Registry Operator is made available.

Instructions are shown for configuring a persistent volume, which is required for production clusters. Where applicable, instructions are shown for configuring an empty directory as the storage location, which is available for only non-production clusters.

Additional instructions are provided for allowing the image registry to use block storage types by using the **Recreate** rollout strategy during upgrades.

15.3.18.3. Configuring block registry storage for bare metal

To allow the image registry to use block storage types during upgrades as a cluster administrator, you can use the **Recreate** rollout strategy.



IMPORTANT

Block storage volumes, or block persistent volumes, are supported but not recommended for use with the image registry on production clusters. An installation where the registry is configured on block storage is not highly available because the registry cannot have more than one replica.

If you choose to use a block storage volume with the image registry, you must use a filesystem persistent volume claim (PVC).

Procedure

1. Enter the following command to set the image registry storage as a block storage type, patch the registry so that it uses the **Recreate** rollout strategy, and runs with only one (**1**) replica:

```
$ oc patch config.imageregistry.operator.openshift.io/cluster --type=merge -p '{"spec": {"rolloutStrategy":"Recreate","replicas":1}}'
```

2. Provision the PV for the block storage device, and create a PVC for that volume. The requested block volume uses the ReadWriteOnce (RWO) access mode.

- a. Create a **pvc.yaml** file with the following contents to define a VMware vSphere **PersistentVolumeClaim** object:

```
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: image-registry-storage 1
  namespace: openshift-image-registry 2
spec:
  accessModes:
    - ReadWriteOnce 3
  resources:
    requests:
      storage: 100Gi 4
```

- 1 A unique name that represents the **PersistentVolumeClaim** object.
- 2 The namespace for the **PersistentVolumeClaim** object, which is **openshift-image-registry**.
- 3 The access mode of the persistent volume claim. With **ReadWriteOnce**, the volume can be mounted with read and write permissions by a single node.
- 4 The size of the persistent volume claim.

- b. Enter the following command to create the **PersistentVolumeClaim** object from the file:

```
$ oc create -f pvc.yaml -n openshift-image-registry
```

3. Enter the following command to edit the registry configuration so that it references the correct PVC:

```
$ oc edit config.imageregistry.operator.openshift.io -o yaml
```

Example output

```
storage:
  pvc:
    claim: 1
```

- 1 By creating a custom PVC, you can leave the **claim** field blank for the default automatic creation of an **image-registry-storage** PVC.

15.3.19. Completing installation on user-provisioned infrastructure

After you complete the Operator configuration, you can finish installing the cluster on infrastructure that you provide.

Prerequisites

- Your control plane has initialized.

- You have completed the initial Operator configuration.

Procedure

- Confirm that all the cluster components are online with the following command:

```
$ watch -n5 oc get clusteroperators
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.16.0	True	False	False	19m
baremetal	4.16.0	True	False	False	37m
cloud-credential	4.16.0	True	False	False	40m
cluster-autoscaler	4.16.0	True	False	False	37m
config-operator	4.16.0	True	False	False	38m
console	4.16.0	True	False	False	26m
csi-snapshot-controller	4.16.0	True	False	False	37m
dns	4.16.0	True	False	False	37m
etcd	4.16.0	True	False	False	36m
image-registry	4.16.0	True	False	False	31m
ingress	4.16.0	True	False	False	30m
insights	4.16.0	True	False	False	31m
kube-apiserver	4.16.0	True	False	False	26m
kube-controller-manager	4.16.0	True	False	False	36m
kube-scheduler	4.16.0	True	False	False	36m
kube-storage-version-migrator	4.16.0	True	False	False	37m
machine-api	4.16.0	True	False	False	29m
machine-approver	4.16.0	True	False	False	37m
machine-config	4.16.0	True	False	False	36m
marketplace	4.16.0	True	False	False	37m
monitoring	4.16.0	True	False	False	29m
network	4.16.0	True	False	False	38m
node-tuning	4.16.0	True	False	False	37m
openshift-apiserver	4.16.0	True	False	False	32m
openshift-controller-manager	4.16.0	True	False	False	30m
openshift-samples	4.16.0	True	False	False	32m
operator-lifecycle-manager	4.16.0	True	False	False	37m
operator-lifecycle-manager-catalog	4.16.0	True	False	False	37m
operator-lifecycle-manager-packageserver	4.16.0	True	False	False	32m
service-ca	4.16.0	True	False	False	38m
storage	4.16.0	True	False	False	37m

Alternatively, the following command notifies you when all of the clusters are available. It also retrieves and displays credentials:

```
$ ./openshift-install --dir <installation_directory> wait-for install-complete 1
```

- For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

Example output

INFO Waiting up to 30m0s for the cluster to initialize...

The command succeeds when the Cluster Version Operator finishes deploying the OpenShift Container Platform cluster from Kubernetes API server.

IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

2. Confirm that the Kubernetes API server is communicating with the pods.

a. To view a list of all pods, use the following command:

```
$ oc get pods --all-namespaces
```

Example output

```

NAMESPACE           NAME                                     READY  STATUS
RESTARTS  AGE
openshift-apiserver-operator  openshift-apiserver-operator-85cb746d55-zqhs8  1/1
Running  1      9m
openshift-apiserver          apiserver-67b9g                                1/1  Running  0
3m
openshift-apiserver          apiserver-ljcmx                                1/1  Running  0
1m
openshift-apiserver          apiserver-z25h4                                1/1  Running  0
2m
openshift-authentication-operator  authentication-operator-69d5d8bf84-vh2n8      1/1
Running  0      5m
...

```

b. View the logs for a pod that is listed in the output of the previous command by using the following command:

```
$ oc logs <pod_name> -n <namespace> 1
```

- 1** Specify the pod name and namespace, as shown in the output of the previous command.

If the pod logs display, the Kubernetes API server can communicate with the cluster machines.

- For an installation with Fibre Channel Protocol (FCP), additional steps are required to enable multipathing. Do not enable multipathing during installation. See "Enabling multipathing with kernel arguments on RHCOS" in the *Postinstallation machine configuration tasks* documentation for more information.

15.3.20. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources


- See [About remote health monitoring](#) for more information about the Telemetry service

15.3.21. Next steps

- [Validating an installation.](#)
- [Customize your cluster.](#)
- If necessary, you can [opt out of remote health reporting](#) .
- [Set up your registry and configure registry storage](#) .

15.4. INSTALLING A USER-PROVISIONED BARE METAL CLUSTER ON A RESTRICTED NETWORK

In OpenShift Container Platform 4.16, you can install a cluster on bare metal infrastructure that you provision in a restricted network.



IMPORTANT

While you might be able to follow this procedure to deploy a cluster on virtualized or cloud environments, you must be aware of additional considerations for non-bare metal platforms. Review the information in the [guidelines for deploying OpenShift Container Platform on non-tested platforms](#) before you attempt to install an OpenShift Container Platform cluster in such an environment.

15.4.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).

- You [created a registry on your mirror host](#) and obtained the `imageContentSources` data for your version of OpenShift Container Platform.



IMPORTANT

Because the installation media is on the mirror host, you can use that computer to complete all installation steps.

- You provisioned [persistent storage](#) for your cluster. To deploy a private image registry, your storage must provide ReadWriteMany access modes.
- If you use a firewall and plan to use the Telemetry service, you [configured the firewall to allow the sites](#) that your cluster requires access to.



NOTE

Be sure to also review this site list if you are configuring a proxy.

15.4.2. About installations in restricted networks

In OpenShift Container Platform 4.16, you can perform an installation that does not require an active connection to the internet to obtain software components. Restricted network installations can be completed using installer-provisioned infrastructure or user-provisioned infrastructure, depending on the cloud platform to which you are installing the cluster.

If you choose to perform a restricted network installation on a cloud platform, you still require access to its cloud APIs. Some cloud functions, like Amazon Web Service's Route 53 DNS and IAM services, require internet access. Depending on your network, you might require less internet access for an installation on bare metal hardware, Nutanix, or on VMware vSphere.

To complete a restricted network installation, you must create a registry that mirrors the contents of the OpenShift image registry and contains the installation media. You can create this registry on a mirror host, which can access both the internet and your closed network, or by using other methods that meet your restrictions.



IMPORTANT

Because of the complexity of the configuration for user-provisioned installations, consider completing a standard user-provisioned infrastructure installation before you attempt a restricted network installation using user-provisioned infrastructure. Completing this test installation might make it easier to isolate and troubleshoot any issues that might arise during your installation in a restricted network.

15.4.2.1. Additional limits

Clusters in restricted networks have the following additional limitations and restrictions:

- The `ClusterVersion` status includes an **Unable to retrieve available updates** error.
- By default, you cannot use the contents of the Developer Catalog because you cannot access the required image stream tags.

15.4.3. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to obtain the images that are necessary to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.

15.4.4. Requirements for a cluster with user-provisioned infrastructure

For a cluster that contains user-provisioned infrastructure, you must deploy all of the required machines.


This section describes the requirements for deploying OpenShift Container Platform on user-provisioned infrastructure.

15.4.4.1. Required machines for cluster installation

The smallest OpenShift Container Platform clusters require the following hosts:

Table 15.32. Minimum required hosts

Hosts	Description
One temporary bootstrap machine	The cluster requires the bootstrap machine to deploy the OpenShift Container Platform cluster on the three control plane machines. You can remove the bootstrap machine after you install the cluster.
Three control plane machines	The control plane machines run the Kubernetes and OpenShift Container Platform services that form the control plane.
At least two compute machines, which are also known as worker machines.	The workloads requested by OpenShift Container Platform users run on the compute machines.



NOTE

As an exception, you can run zero compute machines in a bare metal cluster that consists of three control plane machines only. This provides smaller, more resource efficient clusters for cluster administrators and developers to use for testing, development, and production. Running one compute machine is not supported.



IMPORTANT

To maintain high availability of your cluster, use separate physical hosts for these cluster machines.

The bootstrap and control plane machines must use Red Hat Enterprise Linux CoreOS (RHCOS) as the operating system. However, the compute machines can choose between Red Hat Enterprise Linux CoreOS (RHCOS), Red Hat Enterprise Linux (RHEL) 8.6 and later.

Note that RHCOS is based on Red Hat Enterprise Linux (RHEL) 9.2 and inherits all of its hardware certifications and requirements. See [Red Hat Enterprise Linux technology capabilities and limits](#) .

15.4.4.2. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 15.33. Minimum resource requirements

Machine	Operating System	CPU [1]	RAM	Storage	Input/Output Per Second (IOPS)[2]
Bootstrap	RHCOS	4	16 GB	100 GB	300
Control plane	RHCOS	4	16 GB	100 GB	300
Compute	RHCOS, RHEL 8.6 and later [3]	2	8 GB	100 GB	300

1. One CPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio: $(\text{threads per core} \times \text{cores}) \times \text{sockets} = \text{CPUs}$.
2. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes which require a 10 ms p99 fsync duration. Note that on many cloud platforms, storage size and IOPS scale together, so you might need to over-allocate storage volume to obtain sufficient performance.
3. As with all user-provisioned installations, if you choose to use RHEL compute machines in your cluster, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Use of RHEL 7 compute machines is deprecated and has been removed in OpenShift Container Platform 4.10 and later.



NOTE

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

15.4.4.3. Certificate signing requests management

Because your cluster has limited access to automatic machine management when you use infrastructure that you provision, you must provide a mechanism for approving cluster certificate signing requests (CSRs) after installation. The **kube-controller-manager** only approves the kubelet client CSRs. The **machine-approver** cannot guarantee the validity of a serving certificate that is requested by using kubelet credentials because it cannot confirm that the correct machine issued the request. You must determine and implement a method of verifying the validity of the kubelet serving certificate requests and approving them.

Additional resources

- See [Configuring a three-node cluster](#) for details about deploying three-node clusters in bare metal environments.
- See [Approving the certificate signing requests for your machines](#) for more information about approving cluster certificate signing requests after installation.

15.4.4.4. Networking requirements for user-provisioned infrastructure

All the Red Hat Enterprise Linux CoreOS (RHCOS) machines require networking to be configured in **initramfs** during boot to fetch their Ignition config files.

During the initial boot, the machines require an IP address configuration that is set either through a DHCP server or statically by providing the required boot options. After a network connection is established, the machines download their Ignition config files from an HTTP or HTTPS server. The Ignition config files are then used to set the exact state of each machine. The Machine Config Operator completes more changes to the machines, such as the application of new certificates or keys, after installation.

It is recommended to use a DHCP server for long-term management of the cluster machines. Ensure that the DHCP server is configured to provide persistent IP addresses, DNS server information, and hostnames to the cluster machines.



NOTE

If a DHCP service is not available for your user-provisioned infrastructure, you can instead provide the IP networking configuration and the address of the DNS server to the nodes at RHCOS install time. These can be passed as boot arguments if you are installing from an ISO image. See the *Installing RHCOS and starting the OpenShift Container Platform bootstrap process* section for more information about static IP provisioning and advanced networking options.

The Kubernetes API server must be able to resolve the node names of the cluster machines. If the API servers and worker nodes are in different zones, you can configure a default DNS search zone to allow the API server to resolve the node names. Another supported approach is to always refer to hosts by their fully-qualified domain names in both the node objects and all DNS requests.

15.4.4.4.1. Setting the cluster node hostnames through DHCP

On Red Hat Enterprise Linux CoreOS (RHCOS) machines, the hostname is set through NetworkManager. By default, the machines obtain their hostname through DHCP. If the hostname is not provided by DHCP, set statically through kernel arguments, or another method, it is obtained through a reverse DNS lookup. Reverse DNS lookup occurs after the network has been initialized on a node and can take time to resolve. Other system services can start prior to this and detect the hostname as **localhost** or similar. You can avoid this by using DHCP to provide the hostname for each cluster node.

Additionally, setting the hostnames through DHCP can bypass any manual DNS record name configuration errors in environments that have a DNS split-horizon implementation.

15.4.4.4.2. Network connectivity requirements

You must configure the network connectivity between machines to allow OpenShift Container Platform cluster components to communicate. Each machine must be able to resolve the hostnames of all other machines in the cluster.

This section provides details about the ports that are required.

Table 15.34. Ports used for all-machine to all-machine communications

Protocol	Port	Description
ICMP	N/A	Network reachability tests
TCP	1936	Metrics
	9000-9999	Host level services, including the node exporter on ports 9100-9101 and the Cluster Version Operator on port 9099 .
	10250-10259	The default ports that Kubernetes reserves
UDP	4789	VXLAN

Protocol	Port	Description
	6081	Geneve
	9000-9999	Host level services, including the node exporter on ports 9100-9101 .
	500	IPsec IKE packets
	4500	IPsec NAT-T packets
	123	Network Time Protocol (NTP) on UDP port 123 If an external NTP time server is configured, you must open UDP port 123 .
TCP/UDP	30000-32767	Kubernetes node port
ESP	N/A	IPsec Encapsulating Security Payload (ESP)

Table 15.35. Ports used for all-machine to control plane communications

Protocol	Port	Description
TCP	6443	Kubernetes API

Table 15.36. Ports used for control plane machine to control plane machine communications

Protocol	Port	Description
TCP	2379-2380	etcd server and peer ports

NTP configuration for user-provisioned infrastructure

OpenShift Container Platform clusters are configured to use a public Network Time Protocol (NTP) server by default. If you want to use a local enterprise NTP server, or if your cluster is being deployed in a disconnected network, you can configure the cluster to use a specific time server. For more information, see the documentation for *Configuring chrony time service*.

If a DHCP server provides NTP server information, the chrony time service on the Red Hat Enterprise Linux CoreOS (RHCOS) machines read the information and can sync the clock with the NTP servers.

Additional resources

- [Configuring chrony time service](#)

15.4.4.5. User-provisioned DNS requirements

In OpenShift Container Platform deployments, DNS name resolution is required for the following components:

- The Kubernetes API
- The OpenShift Container Platform application wildcard
- The bootstrap, control plane, and compute machines

Reverse DNS resolution is also required for the Kubernetes API, the bootstrap machine, the control plane machines, and the compute machines.

DNS A/AAAA or CNAME records are used for name resolution and PTR records are used for reverse name resolution. The reverse records are important because Red Hat Enterprise Linux CoreOS (RHCOS) uses the reverse records to set the hostnames for all the nodes, unless the hostnames are provided by DHCP. Additionally, the reverse records are used to generate the certificate signing requests (CSR) that OpenShift Container Platform needs to operate.




NOTE


It is recommended to use a DHCP server to provide the hostnames to each cluster node. See the *DHCP recommendations for user-provisioned infrastructure* section for more information.

The following DNS records are required for a user-provisioned OpenShift Container Platform cluster and they must be in place before installation. In each record, **<cluster_name>** is the cluster name and **<base_domain>** is the base domain that you specify in the **install-config.yaml** file. A complete DNS record takes the form: **<component>.<cluster_name>.<base_domain>..**

Table 15.37. Required DNS records

Component	Record	Description
Kubernetes API	api.<cluster_name>.<base_domain>..	A DNS A/AAAA or CNAME record, and a DNS PTR record, to identify the API load balancer. These records must be resolvable by both clients external to the cluster and from all the nodes within the cluster.
	api-int.<cluster_name>.<base_domain>..	A DNS A/AAAA or CNAME record, and a DNS PTR record, to internally identify the API load balancer. These records must be resolvable from all the nodes within the cluster.
		 <p>IMPORTANT</p> <p>The API server must be able to resolve the worker nodes by the hostnames that are recorded in Kubernetes. If the API server cannot resolve the node names, then proxied API calls can fail, and you cannot retrieve logs from pods.</p>

Component	Record	Description
Routes	*.apps.<cluster_name>.<base_domain>.	<p>A wildcard DNS A/AAAA or CNAME record that refers to the application ingress load balancer. The application ingress load balancer targets the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default. These records must be resolvable by both clients external to the cluster and from all the nodes within the cluster.</p> <p>For example, console-openshift-console.apps.<cluster_name>.<base_domain> is used as a wildcard route to the OpenShift Container Platform console.</p>
Bootstrap machine	bootstrap.<cluster_name>.<base_domain>.	A DNS A/AAAA or CNAME record, and a DNS PTR record, to identify the bootstrap machine. These records must be resolvable by the nodes within the cluster.
Control plane machines	<control_plane><n>.<cluster_name>.<base_domain>.	DNS A/AAAA or CNAME records and DNS PTR records to identify each machine for the control plane nodes. These records must be resolvable by the nodes within the cluster.
Compute machines	<compute><n>.<cluster_name>.<base_domain>.	DNS A/AAAA or CNAME records and DNS PTR records to identify each machine for the worker nodes. These records must be resolvable by the nodes within the cluster.



NOTE

In OpenShift Container Platform 4.4 and later, you do not need to specify etcd host and SRV records in your DNS configuration.

TIP

You can use the **dig** command to verify name and reverse name resolution. See the section on *Validating DNS resolution for user-provisioned infrastructure* for detailed validation steps.

15.4.4.5.1. Example DNS configuration for user-provisioned clusters

This section provides A and PTR record configuration samples that meet the DNS requirements for deploying OpenShift Container Platform on user-provisioned infrastructure. The samples are not meant to provide advice for choosing one DNS solution over another.

In the examples, the cluster name is **ocp4** and the base domain is **example.com**.

Example DNS A record configuration for a user-provisioned cluster

The following example is a BIND zone file that shows sample A records for name resolution in a user-provisioned cluster.

-

Example 15.7. Sample DNS zone database

```

$TTL 1W
@ IN SOA ns1.example.com. root (
  2019070700 ; serial
  3H ; refresh (3 hours)
  30M ; retry (30 minutes)
  2W ; expiry (2 weeks)
  1W) ; minimum (1 week)
IN NS ns1.example.com.
IN MX 10 smtp.example.com.
;
;
ns1.example.com. IN A 192.168.1.5
smtp.example.com. IN A 192.168.1.5
;
helper.example.com. IN A 192.168.1.5
helper.ocp4.example.com. IN A 192.168.1.5
;
api.ocp4.example.com. IN A 192.168.1.5 1
api-int.ocp4.example.com. IN A 192.168.1.5 2
;
*.apps.ocp4.example.com. IN A 192.168.1.5 3
;
bootstrap.ocp4.example.com. IN A 192.168.1.96 4
;
control-plane0.ocp4.example.com. IN A 192.168.1.97 5
control-plane1.ocp4.example.com. IN A 192.168.1.98 6
control-plane2.ocp4.example.com. IN A 192.168.1.99 7
;
compute0.ocp4.example.com. IN A 192.168.1.11 8
compute1.ocp4.example.com. IN A 192.168.1.7 9
;
;EOF

```

- 1 Provides name resolution for the Kubernetes API. The record refers to the IP address of the API load balancer.
- 2 Provides name resolution for the Kubernetes API. The record refers to the IP address of the API load balancer and is used for internal cluster communications.
- 3 Provides name resolution for the wildcard routes. The record refers to the IP address of the application ingress load balancer. The application ingress load balancer targets the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.

**NOTE**

In the example, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

- 4 Provides name resolution for the bootstrap machine.
- 5 6 7 Provides name resolution for the control plane machines.
- 8 9 Provides name resolution for the compute machines.

Example DNS PTR record configuration for a user-provisioned cluster

The following example BIND zone file shows sample PTR records for reverse name resolution in a user-provisioned cluster.

Example 15.8. Sample DNS zone database for reverse records

```
$TTL 1W
@ IN SOA ns1.example.com. root (
    2019070700 ; serial
    3H ; refresh (3 hours)
    30M ; retry (30 minutes)
    2W ; expiry (2 weeks)
    1W ) ; minimum (1 week)
IN NS ns1.example.com.
;
5.1.168.192.in-addr.arpa. IN PTR api.ocp4.example.com. 1
5.1.168.192.in-addr.arpa. IN PTR api-int.ocp4.example.com. 2
;
96.1.168.192.in-addr.arpa. IN PTR bootstrap.ocp4.example.com. 3
;
97.1.168.192.in-addr.arpa. IN PTR control-plane0.ocp4.example.com. 4
98.1.168.192.in-addr.arpa. IN PTR control-plane1.ocp4.example.com. 5
99.1.168.192.in-addr.arpa. IN PTR control-plane2.ocp4.example.com. 6
;
11.1.168.192.in-addr.arpa. IN PTR compute0.ocp4.example.com. 7
7.1.168.192.in-addr.arpa. IN PTR compute1.ocp4.example.com. 8
;
;EOF
```

- 1 Provides reverse DNS resolution for the Kubernetes API. The PTR record refers to the record name of the API load balancer.
- 2 Provides reverse DNS resolution for the Kubernetes API. The PTR record refers to the record name of the API load balancer and is used for internal cluster communications.
- 3 Provides reverse DNS resolution for the bootstrap machine.
- 4 5 6 Provides reverse DNS resolution for the control plane machines.
- 7 8 Provides reverse DNS resolution for the compute machines.

**NOTE**

A PTR record is not required for the OpenShift Container Platform application wildcard.

Additional resources

- [Validating DNS resolution for user-provisioned infrastructure](#)

15.4.4.6. Load balancing requirements for user-provisioned infrastructure

Before you install OpenShift Container Platform, you must provision the API and application Ingress load balancing infrastructure. In production scenarios, you can deploy the API and application Ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

**NOTE**

If you want to deploy the API and application Ingress load balancers with a Red Hat Enterprise Linux (RHEL) instance, you must purchase the RHEL subscription separately.

The load balancing infrastructure must meet the following requirements:

1. **API load balancer.** Provides a common endpoint for users, both human and machine, to interact with and configure the platform. Configure the following conditions:
 - Layer 4 load balancing only. This can be referred to as Raw TCP or SSL Passthrough mode.
 - A stateless load balancing algorithm. The options vary based on the load balancer implementation.

**IMPORTANT**


Do not configure session persistence for an API load balancer. Configuring session persistence for a Kubernetes API server might cause performance issues from excess application traffic for your OpenShift Container Platform cluster and the Kubernetes API that runs inside the cluster.

Configure the following ports on both the front and back of the load balancers:

Table 15.38. API load balancer

Port	Back-end machines (pool members)	Internal	External	Description
6443	Bootstrap and control plane. You remove the bootstrap machine from the load balancer after the bootstrap machine initializes the cluster control plane. You must configure the /readyz endpoint for the API server health check probe.	X	X	Kubernetes API server

Port	Back-end machines (pool members)	Internal	External	Description
22623	Bootstrap and control plane. You remove the bootstrap machine from the load balancer after the bootstrap machine initializes the cluster control plane.	X		Machine config server



NOTE

The load balancer must be configured to take a maximum of 30 seconds from the time the API server turns off the **/readyz** endpoint to the removal of the API server instance from the pool. Within the time frame after **/readyz** returns an error or becomes healthy, the endpoint must have been removed or added. Probing every 5 or 10 seconds, with two successful requests to become healthy and three to become unhealthy, are well-tested values.

2. **Application Ingress load balancer:** Provides an ingress point for application traffic flowing in from outside the cluster. A working configuration for the Ingress router is required for an OpenShift Container Platform cluster.

Configure the following conditions:

- Layer 4 load balancing only. This can be referred to as Raw TCP or SSL Passthrough mode.
- A connection-based or session-based persistence is recommended, based on the options available and types of applications that will be hosted on the platform.

TIP

If the true IP address of the client can be seen by the application Ingress load balancer, enabling source IP-based session persistence can improve performance for applications that use end-to-end TLS encryption.

Configure the following ports on both the front and back of the load balancers:

Table 15.39. Application Ingress load balancer

Port	Back-end machines (pool members)	Internal	External	Description
443	The machines that run the Ingress Controller pods, compute, or worker, by default.	X	X	HTTPS traffic
80	The machines that run the Ingress Controller pods, compute, or worker, by default.	X	X	HTTP traffic

**NOTE**

If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application Ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes.

15.4.4.6.1. Example load balancer configuration for user-provisioned clusters

This section provides an example API and application Ingress load balancer configuration that meets the load balancing requirements for user-provisioned clusters. The sample is an `/etc/haproxy/haproxy.cfg` configuration for an HAProxy load balancer. The example is not meant to provide advice for choosing one load balancing solution over another.

In the example, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

**NOTE**

If you are using HAProxy as a load balancer and SELinux is set to **enforcing**, you must ensure that the HAProxy service can bind to the configured TCP port by running **setsebool -P haproxy_connect_any=1**.

Example 15.9. Sample API and application Ingress load balancer configuration

```

global
  log      127.0.0.1 local2
  pidfile  /var/run/haproxy.pid
  maxconn 4000
  daemon
defaults
  mode          http
  log           global
  option        dontlognull
  option http-server-close
  option        redispatch
  retries       3
  timeout http-request 10s
  timeout queue 1m
  timeout connect 10s
  timeout client 1m
  timeout server 1m
  timeout http-keep-alive 10s
  timeout check 10s
  maxconn      3000
listen api-server-6443 1
  bind *:6443
  mode tcp
  option httpchk GET /readyz HTTP/1.0
  option log-health-checks
  balance roundrobin
  server bootstrap bootstrap.ocp4.example.com:6443 verify none check check-ssl inter 10s fall 2
  rise 3 backup 2
  server master0 master0.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s

```

```

fall 2 rise 3
server master1 master1.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
fall 2 rise 3
server master2 master2.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
fall 2 rise 3
listen machine-config-server-22623 3
bind *:22623
mode tcp
server bootstrap bootstrap.ocp4.example.com:22623 check inter 1s backup 4
server master0 master0.ocp4.example.com:22623 check inter 1s
server master1 master1.ocp4.example.com:22623 check inter 1s
server master2 master2.ocp4.example.com:22623 check inter 1s
listen ingress-router-443 5
bind *:443
mode tcp
balance source
server compute0 compute0.ocp4.example.com:443 check inter 1s
server compute1 compute1.ocp4.example.com:443 check inter 1s
listen ingress-router-80 6
bind *:80
mode tcp
balance source
server compute0 compute0.ocp4.example.com:80 check inter 1s
server compute1 compute1.ocp4.example.com:80 check inter 1s

```

- 1 Port **6443** handles the Kubernetes API traffic and points to the control plane machines.
- 2 4 The bootstrap entries must be in place before the OpenShift Container Platform cluster installation and they must be removed after the bootstrap process is complete.
- 3 Port **22623** handles the machine config server traffic and points to the control plane machines.
- 5 Port **443** handles the HTTPS traffic and points to the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.
- 6 Port **80** handles the HTTP traffic and points to the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.



NOTE

If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application Ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes.

TIP

If you are using HAProxy as a load balancer, you can check that the **haproxy** process is listening on ports **6443**, **22623**, **443**, and **80** by running **netstat -nltpu** on the HAProxy node.

15.4.5. Preparing the user-provisioned infrastructure

Before you install OpenShift Container Platform on user-provisioned infrastructure, you must prepare the underlying infrastructure.

This section provides details about the high-level steps required to set up your cluster infrastructure in preparation for an OpenShift Container Platform installation. This includes configuring IP networking and network connectivity for your cluster nodes, enabling the required ports through your firewall, and setting up the required DNS and load balancing infrastructure.

After preparation, your cluster infrastructure must meet the requirements outlined in the *Requirements for a cluster with user-provisioned infrastructure* section.

Prerequisites

- You have reviewed the [OpenShift Container Platform 4.x Tested Integrations](#) page.
- You have reviewed the infrastructure requirements detailed in the *Requirements for a cluster with user-provisioned infrastructure* section.

Procedure

1. If you are using DHCP to provide the IP networking configuration to your cluster nodes, configure your DHCP service.
 - a. Add persistent IP addresses for the nodes to your DHCP server configuration. In your configuration, match the MAC address of the relevant network interface to the intended IP address for each node.
 - b. When you use DHCP to configure IP addressing for the cluster machines, the machines also obtain the DNS server information through DHCP. Define the persistent DNS server address that is used by the cluster nodes through your DHCP server configuration.



NOTE

If you are not using a DHCP service, you must provide the IP networking configuration and the address of the DNS server to the nodes at RHCOS install time. These can be passed as boot arguments if you are installing from an ISO image. See the *Installing RHCOS and starting the OpenShift Container Platform bootstrap process* section for more information about static IP provisioning and advanced networking options.

- c. Define the hostnames of your cluster nodes in your DHCP server configuration. See the *Setting the cluster node hostnames through DHCP* section for details about hostname considerations.



NOTE

If you are not using a DHCP service, the cluster nodes obtain their hostname through a reverse DNS lookup.

2. Ensure that your network infrastructure provides the required network connectivity between the cluster components. See the *Networking requirements for user-provisioned infrastructure* section for details about the requirements.

3. Configure your firewall to enable the ports required for the OpenShift Container Platform cluster components to communicate. See *Networking requirements for user-provisioned infrastructure* section for details about the ports that are required.



IMPORTANT

By default, port **1936** is accessible for an OpenShift Container Platform cluster, because each control plane node needs access to this port.

Avoid using the Ingress load balancer to expose this port, because doing so might result in the exposure of sensitive information, such as statistics and metrics, related to Ingress Controllers.

4. Setup the required DNS infrastructure for your cluster.
 - a. Configure DNS name resolution for the Kubernetes API, the application wildcard, the bootstrap machine, the control plane machines, and the compute machines.
 - b. Configure reverse DNS resolution for the Kubernetes API, the bootstrap machine, the control plane machines, and the compute machines.
See the *User-provisioned DNS requirements* section for more information about the OpenShift Container Platform DNS requirements.
5. Validate your DNS configuration.
 - a. From your installation node, run DNS lookups against the record names of the Kubernetes API, the wildcard routes, and the cluster nodes. Validate that the IP addresses in the responses correspond to the correct components.
 - b. From your installation node, run reverse DNS lookups against the IP addresses of the load balancer and the cluster nodes. Validate that the record names in the responses correspond to the correct components.
See the *Validating DNS resolution for user-provisioned infrastructure* section for detailed DNS validation steps.
6. Provision the required API and application ingress load balancing infrastructure. See the *Load balancing requirements for user-provisioned infrastructure* section for more information about the requirements.



NOTE

Some load balancing solutions require the DNS name resolution for the cluster nodes to be in place before the load balancing is initialized.


Additional resources

- [Requirements for a cluster with user-provisioned infrastructure](#)
- [Installing RHCOS and starting the OpenShift Container Platform bootstrap process](#)
- [Setting the cluster node hostnames through DHCP](#)
- [Advanced RHCOS installation configuration](#)
- [Networking requirements for user-provisioned infrastructure](#)

- [User-provisioned DNS requirements](#)
- [Validating DNS resolution for user-provisioned infrastructure](#)
- [Load balancing requirements for user-provisioned infrastructure](#)

15.4.6. Validating DNS resolution for user-provisioned infrastructure

You can validate your DNS configuration before installing OpenShift Container Platform on user-provisioned infrastructure.



IMPORTANT

The validation steps detailed in this section must succeed before you install your cluster.

Prerequisites

- You have configured the required DNS records for your user-provisioned infrastructure.

Procedure

1. From your installation node, run DNS lookups against the record names of the Kubernetes API, the wildcard routes, and the cluster nodes. Validate that the IP addresses contained in the responses correspond to the correct components.
 - a. Perform a lookup against the Kubernetes API record name. Check that the result points to the IP address of the API load balancer:

```
$ dig +noall +answer @<nameserver_ip> api.<cluster_name>.<base_domain> 1
```

- 1 Replace **<nameserver_ip>** with the IP address of the nameserver, **<cluster_name>** with your cluster name, and **<base_domain>** with your base domain name.

Example output

```
api.ocp4.example.com. 604800 IN A 192.168.1.5
```

- b. Perform a lookup against the Kubernetes internal API record name. Check that the result points to the IP address of the API load balancer:

```
$ dig +noall +answer @<nameserver_ip> api-int.<cluster_name>.<base_domain>
```

Example output

```
api-int.ocp4.example.com. 604800 IN A 192.168.1.5
```

- c. Test an example ***.apps.<cluster_name>.<base_domain>** DNS wildcard lookup. All of the application wildcard lookups must resolve to the IP address of the application ingress load balancer:

```
$ dig +noall +answer @<nameserver_ip> random.apps.<cluster_name>.<base_domain>
```


Example output

```
random.apps.ocp4.example.com. 604800 IN A 192.168.1.5
```



NOTE

In the example outputs, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

You can replace **random** with another wildcard value. For example, you can query the route to the OpenShift Container Platform console:

```
$ dig +noall +answer @<nameserver_ip> console-openshift-console.apps.
<cluster_name>.<base_domain>
```

Example output

```
console-openshift-console.apps.ocp4.example.com. 604800 IN A 192.168.1.5
```

- d. Run a lookup against the bootstrap DNS record name. Check that the result points to the IP address of the bootstrap node:

```
$ dig +noall +answer @<nameserver_ip> bootstrap.<cluster_name>.<base_domain>
```

Example output

```
bootstrap.ocp4.example.com. 604800 IN A 192.168.1.96
```

- e. Use this method to perform lookups against the DNS record names for the control plane and compute nodes. Check that the results correspond to the IP addresses of each node.
2. From your installation node, run reverse DNS lookups against the IP addresses of the load balancer and the cluster nodes. Validate that the record names contained in the responses correspond to the correct components.
 - a. Perform a reverse lookup against the IP address of the API load balancer. Check that the response includes the record names for the Kubernetes API and the Kubernetes internal API:

```
$ dig +noall +answer @<nameserver_ip> -x 192.168.1.5
```

Example output

```
5.1.168.192.in-addr.arpa. 604800 IN PTR api-int.ocp4.example.com. 1
5.1.168.192.in-addr.arpa. 604800 IN PTR api.ocp4.example.com. 2
```

1 Provides the record name for the Kubernetes internal API.

2 Provides the record name for the Kubernetes API.

**NOTE**

A PTR record is not required for the OpenShift Container Platform application wildcard. No validation step is needed for reverse DNS resolution against the IP address of the application ingress load balancer.

- b. Perform a reverse lookup against the IP address of the bootstrap node. Check that the result points to the DNS record name of the bootstrap node:

```
$ dig +noall +answer @<nameserver_ip> -x 192.168.1.96
```

Example output

```
96.1.168.192.in-addr.arpa. 604800 IN PTR bootstrap.ocp4.example.com.
```

- c. Use this method to perform reverse lookups against the IP addresses for the control plane and compute nodes. Check that the results correspond to the DNS record names of each node.

Additional resources


- [User-provisioned DNS requirements](#)
- [Load balancing requirements for user-provisioned infrastructure](#)

15.4.7. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.

**IMPORTANT**

Do not skip this procedure in production environments, where disaster recovery and debugging is required.

**NOTE**

You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.



NOTE

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.



NOTE

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```



NOTE

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1 Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program. If you install a cluster on infrastructure that you provision, you must provide the key to the installation program.

Additional resources

- [Verifying node health](#)

15.4.8. Manually creating the installation configuration file

Installing the cluster requires that you manually create the installation configuration file.

Prerequisites

- You have an SSH public key on your local machine to provide to the installation program. The key will be used for SSH authentication onto your cluster nodes for debugging and disaster recovery.
- You have obtained the OpenShift Container Platform installation program and the pull secret for your cluster.
- Obtain the **imageContentSources** section from the output of the command to mirror the repository.
- Obtain the contents of the certificate for your mirror registry.

Procedure

1. Create an installation directory to store your required installation assets in:

```
$ mkdir <installation_directory>
```



IMPORTANT

You must create a directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

2. Customize the sample **install-config.yaml** file template that is provided and save it in the **<installation_directory>**.



NOTE

You must name this configuration file **install-config.yaml**.


- Unless you use a registry that RHCOS trusts by default, such as **docker.io**, you must provide the contents of the certificate for your mirror repository in the **additionalTrustBundle** section. In most cases, you must provide the certificate for your mirror.
- You must include the **imageContentSources** section from the output of the command to mirror the repository.



IMPORTANT

- The **ImageContentSourcePolicy** file is generated as an output of **oc mirror** after the mirroring process is finished.
- The **oc mirror** command generates an **ImageContentSourcePolicy** file which contains the YAML needed to define **ImageContentSourcePolicy**. Copy the text from this file and paste it into your **install-config.yaml** file.
- You must run the 'oc mirror' command twice. The first time you run the **oc mirror** command, you get a full **ImageContentSourcePolicy** file. The second time you run the **oc mirror** command, you only get the difference between the first and second run. Because of this behavior, you must always keep a backup of these files in case you need to merge them into one complete **ImageContentSourcePolicy** file. Keeping a backup of these two output files ensures that you have a complete **ImageContentSourcePolicy** file.

3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the next step of the installation process. You must back it up now.

Additional resources

- [Installation configuration parameters for bare metal](#)

15.4.8.1. Sample install-config.yaml file for bare metal

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.

```
apiVersion: v1
baseDomain: example.com 1
compute: 2
- hyperthreading: Enabled 3
name: worker
replicas: 0 4
```

```

controlPlane: 5
  hyperthreading: Enabled 6
  name: master
  replicas: 3 7
metadata:
  name: test 8
networking:
  clusterNetwork:
  - cidr: 10.128.0.0/14 9
    hostPrefix: 23 10
  networkType: OVNKubernetes 11
  serviceNetwork: 12
  - 172.30.0.0/16
platform:
  none: {} 13
fips: false 14
pullSecret: '{"auths":{"<local_registry>":{"auth": "<credentials>","email": "you@example.com"}}}' 15
sshKey: 'ssh-ed25519 AAAA...' 16
additionalTrustBundle: | 17
  -----BEGIN CERTIFICATE-----
  ///////////////////////////////////////////////////////////////////////////////////////////////////////////////////
  -----END CERTIFICATE-----
imageContentSources: 18
- mirrors:
  - <local_registry>/<local_repository_name>/release
  source: quay.io/openshift-release-dev/ocp-release
- mirrors:
  - <local_registry>/<local_repository_name>/release
  source: quay.io/openshift-release-dev/ocp-v4.0-art-dev

```

- 1 The base domain of the cluster. All DNS records must be sub-domains of this base and include the cluster name.
- 2 5 The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.
- 3 6 Specifies whether to enable or disable simultaneous multithreading (SMT), or hyperthreading. By default, SMT is enabled to increase the performance of the cores in your machines. You can disable it by setting the parameter value to **Disabled**. If you disable SMT, you must disable it in all cluster machines; this includes both control plane and compute machines.



NOTE


Simultaneous multithreading (SMT) is enabled by default. If SMT is not enabled in your BIOS settings, the **hyperthreading** parameter has no effect.



IMPORTANT

If you disable **hyperthreading**, whether in the BIOS or in the **install-config.yaml** file, ensure that your capacity planning accounts for the dramatically decreased machine performance.


- 4 You must set this value to **0** when you install OpenShift Container Platform on user-provisioned infrastructure. In installer-provisioned installations, the parameter controls the number of compute



NOTE

If you are installing a three-node cluster, do not deploy any compute machines when you install the Red Hat Enterprise Linux CoreOS (RHCOS) machines.

- 7 The number of control plane machines that you add to the cluster. Because the cluster uses these values as the number of etcd endpoints in the cluster, the value must match the number of control plane machines that you deploy.
- 8 The cluster name that you specified in your DNS records.
- 9 A block of IP addresses from which pod IP addresses are allocated. This block must not overlap with existing physical networks. These IP addresses are used for the pod network. If you need to access the pods from an external network, you must configure load balancers and routers to manage the traffic.



NOTE

Class E CIDR range is reserved for a future use. To use the Class E CIDR range, you must ensure your networking environment accepts the IP addresses within the Class E CIDR range.

- 10 The subnet prefix length to assign to each individual node. For example, if **hostPrefix** is set to **23**, then each node is assigned a **/23** subnet out of the given **cidr**, which allows for 510 ($2^{(32 - 23)} - 2$) pod IP addresses. If you are required to provide access to nodes from an external network, configure load balancers and routers to manage the traffic.
- 11 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 12 The IP address pool to use for service IP addresses. You can enter only one IP address pool. This block must not overlap with existing physical networks. If you need to access the services from an external network, configure load balancers and routers to manage the traffic.
- 13 You must set the platform to **none**. You cannot provide additional platform configuration variables for your platform.



IMPORTANT

Clusters that are installed with the platform type **none** are unable to use some features, such as managing compute machines with the Machine API. This limitation applies even if the compute machines that are attached to the cluster are installed on a platform that would normally support the feature. This parameter cannot be changed after installation.

- 14 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.



IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 15 For **<local_registry>**, specify the registry domain name, and optionally the port, that your mirror registry uses to serve content. For example, **registry.example.com** or **registry.example.com:5000**. For **<credentials>**, specify the base64-encoded user name and password for your mirror registry.
- 16 The SSH public key for the **core** user in Red Hat Enterprise Linux CoreOS (RHCOS).



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- 17 Provide the contents of the certificate file that you used for your mirror registry.
- 18 Provide the **imageContentSources** section according to the output of the command that you used to mirror the repository.



IMPORTANT

- When using the **oc adm release mirror** command, use the output from the **imageContentSources** section.
- When using **oc mirror** command, use the **repositoryDigestMirrors** section of the **ImageContentSourcePolicy** file that results from running the command.
- **ImageContentSourcePolicy** is deprecated. For more information see *Configuring image registry repository mirroring*.

Additional resources

- See [Load balancing requirements for user-provisioned infrastructure](#) for more information on the API and application ingress load balancing requirements.

15.4.8.2. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.



NOTE

For bare metal installations, if you do not assign node IP addresses from the range that is specified in the `networking.machineNetwork[].cidr` field in the `install-config.yaml` file, you must include them in the `proxy.noProxy` field.

Prerequisites

- You have an existing `install-config.yaml` file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the `Proxy` object's `spec.noProxy` field to bypass the proxy if necessary.



NOTE

The `Proxy` object `status.noProxy` field is populated with the values of the `networking.machineNetwork[].cidr`, `networking.clusterNetwork[].cidr`, and `networking.serviceNetwork[]` fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the `Proxy` object `status.noProxy` field is also populated with the instance metadata endpoint (`169.254.169.254`).

Procedure

1. Edit your `install-config.yaml` file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be `http`.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with `.` to match subdomains only. For example, `.y.com` matches `x.y.com`, but not `y.com`. Use `*` to bypass the proxy for all destinations.
- 4 If provided, the installation program generates a config map that is named `user-ca-bundle` in the `openshift-config` namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then

creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.

- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.



NOTE

The installation program does not support the proxy **readinessEndpoints** field.



NOTE

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.



NOTE

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

15.4.8.3. Configuring a three-node cluster

Optionally, you can deploy zero compute machines in a bare metal cluster that consists of three control plane machines only. This provides smaller, more resource efficient clusters for cluster administrators and developers to use for testing, development, and production.

In three-node OpenShift Container Platform environments, the three control plane machines are schedulable, which means that your application workloads are scheduled to run on them.

Prerequisites

- You have an existing **install-config.yaml** file.

Procedure

- Ensure that the number of compute replicas is set to **0** in your **install-config.yaml** file, as shown in the following **compute** stanza:

```
compute:
```

```
- name: worker
  platform: {}
  replicas: 0
```



NOTE

You must set the value of the **replicas** parameter for the compute machines to **0** when you install OpenShift Container Platform on user-provisioned infrastructure, regardless of the number of compute machines you are deploying. In installer-provisioned installations, the parameter controls the number of compute machines that the cluster creates and manages for you. This does not apply to user-provisioned installations, where the compute machines are deployed manually.

For three-node cluster installations, follow these next steps:

- If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes. See the *Load balancing requirements for user-provisioned infrastructure* section for more information.
- When you create the Kubernetes manifest files in the following procedure, ensure that the **mastersSchedulable** parameter in the `<installation_directory>/manifests/cluster-scheduler-02-config.yml` file is set to **true**. This enables your application workloads to run on the control plane nodes.
- Do not deploy any compute nodes when you create the Red Hat Enterprise Linux CoreOS (RHCOS) machines.

15.4.9. Creating the Kubernetes manifest and Ignition config files

Because you must modify some cluster definition files and manually start the cluster machines, you must generate the Kubernetes manifest and Ignition config files that the cluster needs to configure the machines.

The installation configuration file transforms into the Kubernetes manifests. The manifests wrap into the Ignition configuration files, which are later used to configure the cluster machines.



IMPORTANT

- The Ignition config files that the OpenShift Container Platform installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

Prerequisites

- You obtained the OpenShift Container Platform installation program. For a restricted network installation, these files are on your mirror host.
- You created the **install-config.yaml** installation configuration file.

Procedure

1. Change to the directory that contains the OpenShift Container Platform installation program and generate the Kubernetes manifests for the cluster:


```
└─ $ ./openshift-install create manifests --dir <installation_directory> 1
```

- 1 For **<installation_directory>**, specify the installation directory that contains the **install-config.yaml** file you created.



WARNING

If you are installing a three-node cluster, skip the following step to allow the control plane nodes to be schedulable.



IMPORTANT


When you configure control plane nodes from the default unschedulable to schedulable, additional subscriptions are required. This is because control plane nodes then become compute nodes.

2. Check that the **mastersSchedulable** parameter in the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** Kubernetes manifest file is set to **false**. This setting prevents pods from being scheduled on the control plane machines:
 - a. Open the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** file.
 - b. Locate the **mastersSchedulable** parameter and ensure that it is set to **false**.
 - c. Save and exit the file.
3. To create the Ignition configuration files, run the following command from the directory that contains the installation program:

```
└─ $ ./openshift-install create ignition-configs --dir <installation_directory> 1
```

- 1 For **<installation_directory>**, specify the same installation directory.

Ignition config files are created for the bootstrap, control plane, and compute nodes in the installation directory. The **kubeadmin-password** and **kubeconfig** files are created in the `./<installation_directory>/auth` directory:



Additional resources

- See [Recovering from expired control plane certificates](#) for more information about recovering kubelet certificates.

15.4.10. Configuring chrony time service

You must set the time server and related settings used by the chrony time service (**chronyd**) by modifying the contents of the **chrony.conf** file and passing those contents to your nodes as a machine config.

Procedure

1. Create a Butane config including the contents of the **chrony.conf** file. For example, to configure chrony on worker nodes, create a **99-worker-chrony.bu** file.



NOTE

See "Creating machine configs with Butane" for information about Butane.

```

variant: openshift
version: 4.16.0
metadata:
  name: 99-worker-chrony 1
  labels:
    machineconfiguration.openshift.io/role: worker 2
storage:
  files:
  - path: /etc/chrony.conf
    mode: 0644 3
    overwrite: true
    contents:
      inline: |
        pool 0.rhel.pool.ntp.org iburst 4
        driftfile /var/lib/chrony/drift
        makestep 1.0 3
        rtsync
        logdir /var/log/chrony

```

- 1** **2** On control plane nodes, substitute **master** for **worker** in both of these locations.

- 3 Specify an octal value mode for the **mode** field in the machine config file. After creating the file and applying the changes, the **mode** is converted to a decimal value. You can check
 - 4 Specify any valid, reachable time source, such as the one provided by your DHCP server.
2. Use Butane to generate a **MachineConfig** object file, **99-worker-chrony.yaml**, containing the configuration to be delivered to the nodes:

```
$ butane 99-worker-chrony.bu -o 99-worker-chrony.yaml
```

3. Apply the configurations in one of two ways:
 - If the cluster is not running yet, after you generate manifest files, add the **MachineConfig** object file to the **<installation_directory>/openshift** directory, and then continue to create the cluster.
 - If the cluster is already running, apply the file:

```
$ oc apply -f ./99-worker-chrony.yaml
```

15.4.11. Installing RHCOS and starting the OpenShift Container Platform bootstrap process

To install OpenShift Container Platform on bare metal infrastructure that you provision, you must install Red Hat Enterprise Linux CoreOS (RHCOS) on the machines. When you install RHCOS, you must provide the Ignition config file that was generated by the OpenShift Container Platform installation program for the type of machine you are installing. If you have configured suitable networking, DNS, and load balancing infrastructure, the OpenShift Container Platform bootstrap process begins automatically after the RHCOS machines have rebooted.

To install RHCOS on the machines, follow either the steps to use an ISO image or network PXE booting.



NOTE

The compute node deployment steps included in this installation document are RHCOS-specific. If you choose instead to deploy RHEL-based compute nodes, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Only RHEL 8 compute machines are supported.

You can configure RHCOS during ISO and PXE installations by using the following methods:

- **Kernel arguments:** You can use kernel arguments to provide installation-specific information. For example, you can specify the locations of the RHCOS installation files that you uploaded to your HTTP server and the location of the Ignition config file for the type of node you are installing. For a PXE installation, you can use the **APPEND** parameter to pass the arguments to the kernel of the live installer. For an ISO installation, you can interrupt the live installation boot process to add the kernel arguments. In both installation cases, you can use special **coreos.inst.*** arguments to direct the live installer, as well as standard installation boot arguments for turning standard kernel services on or off.
- **Ignition configs:** OpenShift Container Platform Ignition config files (***.ign**) are specific to the type of node you are installing. You pass the location of a bootstrap, control plane, or compute node Ignition config file during the RHCOS installation so that it takes effect on first boot. In

special cases, you can create a separate, limited Ignition config to pass to the live system. That Ignition config could do a certain set of tasks, such as reporting success to a provisioning system after completing installation. This special Ignition config is consumed by the **coreos-installer** to be applied on first boot of the installed system. Do not provide the standard control plane and compute node Ignition configs to the live ISO directly.

- **coreos-installer:** You can boot the live ISO installer to a shell prompt, which allows you to prepare the permanent system in a variety of ways before first boot. In particular, you can run the **coreos-installer** command to identify various artifacts to include, work with disk partitions, and set up networking. In some cases, you can configure features on the live system and copy them to the installed system.

Whether to use an ISO or PXE install depends on your situation. A PXE install requires an available DHCP service and more preparation, but can make the installation process more automated. An ISO install is a more manual process and can be inconvenient if you are setting up more than a few machines.



NOTE

As of OpenShift Container Platform 4.6, the RHCOS ISO and other installation artifacts provide support for installation on disks with 4K sectors.

15.4.11.1. Installing RHCOS by using an ISO image

You can use an ISO image to install RHCOS on the machines.

Prerequisites

- You have created the Ignition config files for your cluster.
- You have configured suitable network, DNS and load balancing infrastructure.
- You have an HTTP server that can be accessed from your computer, and from the machines that you create.
- You have reviewed the *Advanced RHCOS installation configuration* section for different ways to configure features, such as networking and disk partitioning.


Procedure

1. Obtain the SHA512 digest for each of your Ignition config files. For example, you can use the following on a system running Linux to get the SHA512 digest for your **bootstrap.ign** Ignition config file:

```
$ sha512sum <installation_directory>/bootstrap.ign
```

The digests are provided to the **coreos-installer** in a later step to validate the authenticity of the Ignition config files on the cluster nodes.

2. Upload the bootstrap, control plane, and compute node Ignition config files that the installation program created to your HTTP server. Note the URLs of these files.



IMPORTANT

You can add or change configuration settings in your Ignition configs before saving them to your HTTP server. If you plan to add more compute machines to your cluster after you finish installation, do not delete these files.

- From the installation host, validate that the Ignition config files are available on the URLs. The following example gets the Ignition config file for the bootstrap node:

```
$ curl -k http://<HTTP_server>/bootstrap.ign 1
```

Example output

```
% Total    % Received % Xferd Average Speed   Time    Time     Time  Current
           Dload  Upload  Total   Spent    Left    Speed
  0   0   0   0   0   0   0   0   0  --:--:-- --:--:-- --:--:--    0{"ignition":
{"version":"3.2.0"},"passwd":{"users":[{"name":"core","sshAuthorizedKeys":["ssh-rsa...
```


Replace **bootstrap.ign** with **master.ign** or **worker.ign** in the command to validate that the Ignition config files for the control plane and compute nodes are also available.

- Although it is possible to obtain the RHCOS images that are required for your preferred method of installing operating system instances from the [RHCOS image mirror](#) page, the recommended way to obtain the correct version of your RHCOS images are from the output of **openshift-install** command:

```
$ openshift-install coreos print-stream-json | grep '\.iso[^\.]'
```

Example output

```
"location": "<url>/art/storage/releases/rhcos-4.16-aarch64/<release>/aarch64/rhcos-
<release>-live.aarch64.iso",
"location": "<url>/art/storage/releases/rhcos-4.16-ppc64le/<release>/ppc64le/rhcos-
<release>-live.ppc64le.iso",
"location": "<url>/art/storage/releases/rhcos-4.16-s390x/<release>/s390x/rhcos-<release>-
live.s390x.iso",
"location": "<url>/art/storage/releases/rhcos-4.16/<release>/x86_64/rhcos-<release>-
live.x86_64.iso",
```



IMPORTANT

The RHCOS images might not change with every release of OpenShift Container Platform. You must download images with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Use the image versions that match your OpenShift Container Platform version if they are available. Use only ISO images for this procedure. RHCOS qcow2 images are not supported for this installation type.

ISO file names resemble the following example:

rhcos-<version>-live.<architecture>.iso

- Use the ISO to start the RHCOS installation. Use one of the following installation options:

- Burn the ISO image to a disk and boot it directly.
 - Use ISO redirection by using a lights-out management (LOM) interface.
6. Boot the RHCOS ISO image without specifying any options or interrupting the live boot sequence. Wait for the installer to boot into a shell prompt in the RHCOS live environment.



NOTE

It is possible to interrupt the RHCOS installation boot process to add kernel arguments. However, for this ISO procedure you should use the **coreos-installer** command as outlined in the following steps, instead of adding kernel arguments.

7. Run the **coreos-installer** command and specify the options that meet your installation requirements. At a minimum, you must specify the URL that points to the Ignition config file for the node type, and the device that you are installing to:

```
$ sudo coreos-installer install --ignition-url=http://<HTTP_server>/<node_type>.ign <device>
--ignition-hash=sha512-<digest> 1 2
```

- 1 1 You must run the **coreos-installer** command by using **sudo**, because the **core** user does not have the required root privileges to perform the installation.

- 2 The **--ignition-hash** option is required when the Ignition config file is obtained through an HTTP URL to validate the authenticity of the Ignition config file on the cluster node. **<digest>** is the Ignition config file SHA512 digest obtained in a preceding step.




NOTE

If you want to provide your Ignition config files through an HTTPS server that uses TLS, you can add the internal certificate authority (CA) to the system trust store before running **coreos-installer**.

The following example initializes a bootstrap node installation to the **/dev/sda** device. The Ignition config file for the bootstrap node is obtained from an HTTP web server with the IP address 192.168.1.2:

```
$ sudo coreos-installer install --ignition-
url=http://192.168.1.2:80/installation_directory/bootstrap.ign /dev/sda --ignition-hash=sha512-
a5a2d43879223273c9b60af66b44202a1d1248fc01cf156c46d4a79f552b6bad47bc8cc78ddf011
6e80c59d2ea9e32ba53bc807afbca581aa059311def2c3e3b
```

8. Monitor the progress of the RHCOS installation on the console of the machine.



IMPORTANT

Be sure that the installation is successful on each node before commencing with the OpenShift Container Platform installation. Observing the installation process can also help to determine the cause of RHCOS installation issues that might arise.

9. After RHCOS installs, you must reboot the system. During the system reboot, it applies the Ignition config file that you specified.
10. Check the console output to verify that Ignition ran.

Example command

```
Ignition: ran on 2022/03/14 14:48:33 UTC (this boot)
Ignition: user-provided config was applied
```

11. Continue to create the other machines for your cluster.



IMPORTANT

You must create the bootstrap and control plane machines at this time. If the control plane machines are not made schedulable, also create at least two compute machines before you install OpenShift Container Platform.

If the required network, DNS, and load balancer infrastructure are in place, the OpenShift Container Platform bootstrap process begins automatically after the RHCOS nodes have rebooted.



NOTE

RHCOS nodes do not include a default password for the **core** user. You can access the nodes by running **ssh core@<node>.<cluster_name>.<base_domain>** as a user with access to the SSH private key that is paired to the public key that you specified in your **install_config.yaml** file. OpenShift Container Platform 4 cluster nodes running RHCOS are immutable and rely on Operators to apply cluster changes. Accessing cluster nodes by using SSH is not recommended. However, when investigating installation issues, if the OpenShift Container Platform API is not available, or the kubelet is not properly functioning on a target node, SSH access might be required for debugging or disaster recovery.

15.4.11.2. Installing RHCOS by using PXE or iPXE booting


You can use PXE or iPXE booting to install RHCOS on the machines.

Prerequisites

- You have created the Ignition config files for your cluster.
- You have configured suitable network, DNS and load balancing infrastructure.
- You have configured suitable PXE or iPXE infrastructure.
- You have an HTTP server that can be accessed from your computer, and from the machines that you create.
- You have reviewed the *Advanced RHCOS installation configuration* section for different ways to configure features, such as networking and disk partitioning.

Procedure

1. Upload the bootstrap, control plane, and compute node Ignition config files that the installation program created to your HTTP server. Note the URLs of these files.



IMPORTANT

You can add or change configuration settings in your Ignition configs before saving them to your HTTP server. If you plan to add more compute machines to your cluster after you finish installation, do not delete these files.

2. From the installation host, validate that the Ignition config files are available on the URLs. The following example gets the Ignition config file for the bootstrap node:

```
$ curl -k http://<HTTP_server>/bootstrap.ign ❶
```

Example output

```
% Total % Received % Xferd Average Speed Time Time Time Current
          Dload Upload Total Spent Left Speed
  0  0  0  0  0  0  0  0 --:--:-- --:--:-- --:--:-- 0{"ignition":
{"version":"3.2.0"},"passwd":{"users":[{"name":"core","sshAuthorizedKeys":["ssh-rsa...
```

Replace **bootstrap.ign** with **master.ign** or **worker.ign** in the command to validate that the Ignition config files for the control plane and compute nodes are also available.


3. Although it is possible to obtain the RHCOS **kernel**, **initramfs** and **rootfs** files that are required for your preferred method of installing operating system instances from the [RHCOS image mirror](#) page, the recommended way to obtain the correct version of your RHCOS files are from the output of **openshift-install** command:

```
$ openshift-install coreos print-stream-json | grep -Eo "https.*(kernel-|initramfs-|rootfs-)\w+(\.img)?"
```

Example output

```
"<url>/art/storage/releases/rhcos-4.16-aarch64/<release>/aarch64/rhcos-<release>-live-
kernel-aarch64"
"<url>/art/storage/releases/rhcos-4.16-aarch64/<release>/aarch64/rhcos-<release>-live-
initramfs.aarch64.img"
"<url>/art/storage/releases/rhcos-4.16-aarch64/<release>/aarch64/rhcos-<release>-live-
rootfs.aarch64.img"
"<url>/art/storage/releases/rhcos-4.16-ppc64le/49.84.202110081256-0/ppc64le/rhcos-
<release>-live-kernel-ppc64le"
"<url>/art/storage/releases/rhcos-4.16-ppc64le/<release>/ppc64le/rhcos-<release>-live-
initramfs.ppc64le.img"
"<url>/art/storage/releases/rhcos-4.16-ppc64le/<release>/ppc64le/rhcos-<release>-live-
rootfs.ppc64le.img"
"<url>/art/storage/releases/rhcos-4.16-s390x/<release>/s390x/rhcos-<release>-live-kernel-
s390x"
"<url>/art/storage/releases/rhcos-4.16-s390x/<release>/s390x/rhcos-<release>-live-
initramfs.s390x.img"
"<url>/art/storage/releases/rhcos-4.16-s390x/<release>/s390x/rhcos-<release>-live-
rootfs.s390x.img"
"<url>/art/storage/releases/rhcos-4.16/<release>/x86_64/rhcos-<release>-live-kernel-
```

```
x86_64"
"<url>/art/storage/releases/rhcos-4.16/<release>/x86_64/rhcos-<release>-live-
initramfs.x86_64.img"
"<url>/art/storage/releases/rhcos-4.16/<release>/x86_64/rhcos-<release>-live-
rootfs.x86_64.img"
```




IMPORTANT

The RHCOS artifacts might not change with every release of OpenShift Container Platform. You must download images with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Only use the appropriate **kernel**, **initramfs**, and **rootfs** artifacts described below for this procedure. RHCOS QCOW2 images are not supported for this installation type.

The file names contain the OpenShift Container Platform version number. They resemble the following examples:

- **kernel:** rhcos-<version>-live-kernel-<architecture>
- **initramfs:** rhcos-<version>-live-initramfs.<architecture>.img
- **rootfs:** rhcos-<version>-live-rootfs.<architecture>.img

4. Upload the **rootfs**, **kernel**, and **initramfs** files to your HTTP server.



IMPORTANT

If you plan to add more compute machines to your cluster after you finish installation, do not delete these files.

5. Configure the network boot infrastructure so that the machines boot from their local disks after RHCOS is installed on them.
6. Configure PXE or iPXE installation for the RHCOS images and begin the installation. Modify one of the following example menu entries for your environment and verify that the image and Ignition files are properly accessible:

- For PXE (**x86_64**):

```
DEFAULT pxeboot
TIMEOUT 20
PROMPT 0
LABEL pxeboot
  KERNEL http://<HTTP_server>/rhcos-<version>-live-kernel-<architecture> 1
  APPEND initrd=http://<HTTP_server>/rhcos-<version>-live-initramfs.
<architecture>.img coreos.live.rootfs_url=http://<HTTP_server>/rhcos-<version>-live-
rootfs.<architecture>.img coreos.inst.install_dev=/dev/sda
coreos.inst.ignition_url=http://<HTTP_server>/bootstrap.ign 2 3
```

- 1** Specify the location of the live **kernel** file that you uploaded to your HTTP server. The URL must be HTTP, TFTP, or FTP; HTTPS and NFS are not supported.
- 2** If you use multiple NICs, specify a single interface in the **ip** option. For example, to use DHCP on a NIC that is named **enp1**, set **ip=enp1:dhcp**
- 3**

DHCP on a NIC that is named **eno1**, set **ip=eno1:dhcp**.

- 3 Specify the locations of the RHCOS files that you uploaded to your HTTP server. The **initrd** parameter value is the location of the **initramfs** file, the **coreos.live.rootfs_url** parameter value is the location of the **rootfs** file, and the **coreos.inst.ignition_url** parameter value is the location of the bootstrap Ignition config file. You can also add more kernel arguments to the **APPEND** line to configure networking or other boot options.



NOTE

This configuration does not enable serial console access on machines with a graphical console. To configure a different console, add one or more **console=** arguments to the **APPEND** line. For example, add **console=tty0 console=ttyS0** to set the first PC serial port as the primary console and the graphical console as a secondary console. For more information, see [How does one set up a serial terminal and/or console in Red Hat Enterprise Linux?](#) and "Enabling the serial console for PXE and ISO installation" in the "Advanced RHCOS installation configuration" section.

- For iPXE (**x86_64** + **aarch64**):

```
kernel http://<HTTP_server>/rhcos-<version>-live-kernel-<architecture> initrd=main
coreos.live.rootfs_url=http://<HTTP_server>/rhcos-<version>-live-rootfs.
<architecture>.img coreos.inst.install_dev=/dev/sda
coreos.inst.ignition_url=http://<HTTP_server>/bootstrap.ign 1 2
initrd --name main http://<HTTP_server>/rhcos-<version>-live-initramfs.
<architecture>.img 3
boot
```

- 1 Specify the locations of the RHCOS files that you uploaded to your HTTP server. The **kernel** parameter value is the location of the **kernel** file, the **initrd=main** argument is needed for booting on UEFI systems, the **coreos.live.rootfs_url** parameter value is the location of the **rootfs** file, and the **coreos.inst.ignition_url** parameter value is the location of the bootstrap Ignition config file.
- 2 If you use multiple NICs, specify a single interface in the **ip** option. For example, to use DHCP on a NIC that is named **eno1**, set **ip=eno1:dhcp**.
- 3 Specify the location of the **initramfs** file that you uploaded to your HTTP server.



NOTE

This configuration does not enable serial console access on machines with a graphical console. To configure a different console, add one or more **console=** arguments to the **kernel** line. For example, add **console=tty0 console=ttyS0** to set the first PC serial port as the primary console and the graphical console as a secondary console. For more information, see [How does one set up a serial terminal and/or console in Red Hat Enterprise Linux?](#) and "Enabling the serial console for PXE and ISO installation" in the "Advanced RHCOS installation configuration" section.

**NOTE**

To network boot the CoreOS **kernel** on **aarch64** architecture, you need to use a version of iPXE build with the **IMAGE_GZIP** option enabled. See [IMAGE_GZIP option in iPXE](#).

- For PXE (with UEFI and Grub as second stage) on **aarch64**:

```

menuentry 'Install CoreOS' {
  linux rhcos-<version>-live-kernel-<architecture>
  coreos.live.rootfs_url=http://<HTTP_server>/rhcos-<version>-live-rootfs.
  <architecture>.img coreos.inst.install_dev=/dev/sda
  coreos.inst.ignition_url=http://<HTTP_server>/bootstrap.ign 1 2
  initrd rhcos-<version>-live-initramfs.<architecture>.img 3
}

```

- 1** Specify the locations of the RHCOS files that you uploaded to your HTTP/TFTP server. The **kernel** parameter value is the location of the **kernel** file on your TFTP server. The **coreos.live.rootfs_url** parameter value is the location of the **rootfs** file, and the **coreos.inst.ignition_url** parameter value is the location of the bootstrap Ignition config file on your HTTP Server.
- 2** If you use multiple NICs, specify a single interface in the **ip** option. For example, to use DHCP on a NIC that is named **eno1**, set **ip=eno1:dhcp**.
- 3** Specify the location of the **initramfs** file that you uploaded to your TFTP server.

7. Monitor the progress of the RHCOS installation on the console of the machine.

**IMPORTANT**

Be sure that the installation is successful on each node before commencing with the OpenShift Container Platform installation. Observing the installation process can also help to determine the cause of RHCOS installation issues that might arise.

8. After RHCOS installs, the system reboots. During reboot, the system applies the Ignition config file that you specified.
9. Check the console output to verify that Ignition ran.


Example command

```

Ignition: ran on 2022/03/14 14:48:33 UTC (this boot)
Ignition: user-provided config was applied

```

10. Continue to create the machines for your cluster.

**IMPORTANT**

You must create the bootstrap and control plane machines at this time. If the control plane machines are not made schedulable, also create at least two compute machines before you install the cluster.

If the required network, DNS, and load balancer infrastructure are in place, the OpenShift Container Platform bootstrap process begins automatically after the RHCOS nodes have rebooted.



NOTE

RHCOS nodes do not include a default password for the **core** user. You can access the nodes by running **ssh core@<node>.<cluster_name>.<base_domain>** as a user with access to the SSH private key that is paired to the public key that you specified in your **install_config.yaml** file. OpenShift Container Platform 4 cluster nodes running RHCOS are immutable and rely on Operators to apply cluster changes. Accessing cluster nodes by using SSH is not recommended. However, when investigating installation issues, if the OpenShift Container Platform API is not available, or the kubelet is not properly functioning on a target node, SSH access might be required for debugging or disaster recovery.

15.4.11.3. Advanced RHCOS installation configuration

A key benefit for manually provisioning the Red Hat Enterprise Linux CoreOS (RHCOS) nodes for OpenShift Container Platform is to be able to do configuration that is not available through default OpenShift Container Platform installation methods. This section describes some of the configurations that you can do using techniques that include:

- Passing kernel arguments to the live installer
- Running **coreos-installer** manually from the live system
- Customizing a live ISO or PXE boot image

The advanced configuration topics for manual Red Hat Enterprise Linux CoreOS (RHCOS) installations detailed in this section relate to disk partitioning, networking, and using Ignition configs in different ways.

15.4.11.3.1. Using advanced networking options for PXE and ISO installations

Networking for OpenShift Container Platform nodes uses DHCP by default to gather all necessary configuration settings. To set up static IP addresses or configure special settings, such as bonding, you can do one of the following:

- Pass special kernel parameters when you boot the live installer.
- Use a machine config to copy networking files to the installed system.
- Configure networking from a live installer shell prompt, then copy those settings to the installed system so that they take effect when the installed system first boots.

To configure a PXE or iPXE installation, use one of the following options:

- See the "Advanced RHCOS installation reference" tables.
- Use a machine config to copy networking files to the installed system.

To configure an ISO installation, use the following procedure.

Procedure

1. Boot the ISO installer.
2. From the live system shell prompt, configure networking for the live system using available RHEL tools, such as **nmcli** or **nmtui**.
3. Run the **coreos-installer** command to install the system, adding the **--copy-network** option to copy networking configuration. For example:

```
$ sudo coreos-installer install --copy-network \  
--ignition-url=http://host/worker.ign /dev/disk/by-id/scsi-<serial_number>
```



IMPORTANT

The **--copy-network** option only copies networking configuration found under **/etc/NetworkManager/system-connections**. In particular, it does not copy the system hostname.

4. Reboot into the installed system.

Additional resources

- See [Getting started with nmcli](#) and [Getting started with nmtui](#) in the RHEL 8 documentation for more information about the **nmcli** and **nmtui** tools.

15.4.11.3.2. Disk partitioning

Disk partitions are created on OpenShift Container Platform cluster nodes during the Red Hat Enterprise Linux CoreOS (RHCOS) installation. Each RHCOS node of a particular architecture uses the same partition layout, unless you override the default partitioning configuration. During the RHCOS installation, the size of the root file system is increased to use any remaining available space on the target device.



IMPORTANT


The use of a custom partition scheme on your node might result in OpenShift Container Platform not monitoring or alerting on some node partitions. If you override the default partitioning, see [Understanding OpenShift File System Monitoring \(eviction conditions\)](#) for more information about how OpenShift Container Platform monitors your host file systems.

OpenShift Container Platform monitors the following two filesystem identifiers:

- **nodefs**, which is the filesystem that contains **/var/lib/kubelet**
- **imagefs**, which is the filesystem that contains **/var/lib/containers**

For the default partition scheme, **nodefs** and **imagefs** monitor the same root filesystem, **/**.

To override the default partitioning when installing RHCOS on an OpenShift Container Platform cluster node, you must create separate partitions. Consider a situation where you want to add a separate storage partition for your containers and container images. For example, by mounting **/var/lib/containers** in a separate partition, the kubelet separately monitors **/var/lib/containers** as the **imagefs** directory and the root file system as the **nodefs** directory.



IMPORTANT


If you have resized your disk size to host a larger file system, consider creating a separate **/var/lib/containers** partition. Consider resizing a disk that has an **xfs** format to reduce CPU time issues caused by a high number of allocation groups.

15.4.11.3.2.1. Creating a separate **/var** partition

In general, you should use the default disk partitioning that is created during the RHCOS installation. However, there are cases where you might want to create a separate partition for a directory that you expect to grow.

OpenShift Container Platform supports the addition of a single partition to attach storage to either the **/var** directory or a subdirectory of **/var**. For example:

- **/var/lib/containers**: Holds container-related content that can grow as more images and containers are added to a system.
- **/var/lib/etcd**: Holds data that you might want to keep separate for purposes such as performance optimization of etcd storage.
- **/var**: Holds data that you might want to keep separate for purposes such as auditing.



IMPORTANT

For disk sizes larger than 100GB, and especially larger than 1TB, create a separate **/var** partition.

Storing the contents of a **/var** directory separately makes it easier to grow storage for those areas as needed and reinstall OpenShift Container Platform at a later date and keep that data intact. With this method, you will not have to pull all your containers again, nor will you have to copy massive log files when you update systems.

The use of a separate partition for the **/var** directory or a subdirectory of **/var** also prevents data growth in the partitioned directory from filling up the root file system.

The following procedure sets up a separate **/var** partition by adding a machine config manifest that is wrapped into the Ignition config file for a node type during the preparation phase of an installation.

Procedure

1. On your installation host, change to the directory that contains the OpenShift Container Platform installation program and generate the Kubernetes manifests for the cluster:

```
$ openshift-install create manifests --dir <installation_directory>
```

2. Create a Butane config that configures the additional partition. For example, name the file **\$HOME/clusterconfig/98-var-partition.bu**, change the disk device name to the name of the storage device on the **worker** systems, and set the storage size as appropriate. This example places the **/var** directory on a separate partition:


```
variant: openshift
version: 4.16.0
metadata:
  labels:
```

```

machineconfiguration.openshift.io/role: worker
name: 98-var-partition
storage:
  disks:
    - device: /dev/disk/by-id/<device_name> ❶
      partitions:
        - label: var
          start_mib: <partition_start_offset> ❷
          size_mib: <partition_size> ❸
          number: 5
      filesystems:
        - device: /dev/disk/by-partlabel/var
          path: /var
          format: xfs
          mount_options: [defaults, prjquota] ❹
          with_mount_unit: true

```

- ❶ The storage device name of the disk that you want to partition.
- ❷ When adding a data partition to the boot disk, a minimum offset value of 25000 mebibytes is recommended. The root file system is automatically resized to fill all available space up to the specified offset. If no offset value is specified, or if the specified value is smaller than the recommended minimum, the resulting root file system will be too small, and future reinstalls of RHCOS might overwrite the beginning of the data partition.
- ❸ The size of the data partition in mebibytes.
- ❹ The **prjquota** mount option must be enabled for filesystems used for container storage.



NOTE

When creating a separate **/var** partition, you cannot use different instance types for compute nodes, if the different instance types do not have the same device name.

3. Create a manifest from the Butane config and save it to the **clusterconfig/openshift** directory. For example, run the following command:

```
$ butane $HOME/clusterconfig/98-var-partition.bu -o $HOME/clusterconfig/openshift/98-var-partition.yaml
```


4. Create the Ignition config files:

```
$ openshift-install create ignition-configs --dir <installation_directory> ❶
```

- ❶ For **<installation_directory>**, specify the same installation directory.

Ignition config files are created for the bootstrap, control plane, and compute nodes in the installation directory:

```
├── auth
```



The files in the `<installation_directory>/manifest` and `<installation_directory>/openshift` directories are wrapped into the Ignition config files, including the file that contains the **98-var-partition** custom **MachineConfig** object.

Next steps

- You can apply the custom disk partitioning by referencing the Ignition config files during the RHCOS installations.

15.4.11.3.2.2. Retaining existing partitions

For an ISO installation, you can add options to the **coreos-installer** command that cause the installer to maintain one or more existing partitions. For a PXE installation, you can add **coreos.inst.*** options to the **APPEND** parameter to preserve partitions.

Saved partitions might be data partitions from an existing OpenShift Container Platform system. You can identify the disk partitions you want to keep either by partition label or by number.



NOTE

If you save existing partitions, and those partitions do not leave enough space for RHCOS, the installation will fail without damaging the saved partitions.

Retaining existing partitions during an ISO installation

This example preserves any partition in which the partition label begins with **data** (**data***):

```
# coreos-installer install --ignition-url http://10.0.2.2:8080/user.ign \
  --save-partlabel 'data*' /dev/disk/by-id/scsi-<serial_number>
```

The following example illustrates running the **coreos-installer** in a way that preserves the sixth (6) partition on the disk:

```
# coreos-installer install --ignition-url http://10.0.2.2:8080/user.ign \
  --save-partindex 6 /dev/disk/by-id/scsi-<serial_number>
```

This example preserves partitions 5 and higher:

```
# coreos-installer install --ignition-url http://10.0.2.2:8080/user.ign
  --save-partindex 5- /dev/disk/by-id/scsi-<serial_number>
```

In the previous examples where partition saving is used, **coreos-installer** recreates the partition immediately.

Retaining existing partitions during a PXE installation

This **APPEND** option preserves any partition in which the partition label begins with 'data' ('data*'):

```
coreos.inst.save_partlabel=data*
```

This **APPEND** option preserves partitions 5 and higher:

```
coreos.inst.save_partindex=5-
```

This **APPEND** option preserves partition 6:

```
coreos.inst.save_partindex=6
```

15.4.11.3.3. Identifying Ignition configs

When doing an RHCOS manual installation, there are two types of Ignition configs that you can provide, with different reasons for providing each one:

- **Permanent install Ignition config:** Every manual RHCOS installation needs to pass one of the Ignition config files generated by **openshift-installer**, such as **bootstrap.ign**, **master.ign** and **worker.ign**, to carry out the installation.



IMPORTANT

It is not recommended to modify these Ignition config files directly. You can update the manifest files that are wrapped into the Ignition config files, as outlined in examples in the preceding sections.

For PXE installations, you pass the Ignition configs on the **APPEND** line using the **coreos.inst.ignition_url=** option. For ISO installations, after the ISO boots to the shell prompt, you identify the Ignition config on the **coreos-installer** command line with the **--ignition-url=** option. In both cases, only HTTP and HTTPS protocols are supported.

- **Live install Ignition config:** This type can be created by using the **coreos-installer customize** subcommand and its various options. With this method, the Ignition config passes to the live install medium, runs immediately upon booting, and performs setup tasks before or after the RHCOS system installs to disk. This method should only be used for performing tasks that must be done once and not applied again later, such as with advanced partitioning that cannot be done using a machine config.

For PXE or ISO boots, you can create the Ignition config and **APPEND** the **ignition.config.url=** option to identify the location of the Ignition config. You also need to append **ignition.firstboot** **ignition.platform.id=metal** or the **ignition.config.url** option will be ignored.

15.4.11.3.4. Default console configuration

Red Hat Enterprise Linux CoreOS (RHCOS) nodes installed from an OpenShift Container Platform 4.16 boot image use a default console that is meant to accommodate most virtualized and bare metal setups. Different cloud and virtualization platforms may use different default settings depending on the chosen architecture. Bare metal installations use the kernel default settings which typically means the graphical console is the primary console and the serial console is disabled.

The default consoles may not match your specific hardware configuration or you might have specific needs that require you to adjust the default console. For example:

- You want to access the emergency shell on the console for debugging purposes.

- Your cloud platform does not provide interactive access to the graphical console, but provides a serial console.
- You want to enable multiple consoles.

Console configuration is inherited from the boot image. This means that new nodes in existing clusters are unaffected by changes to the default console.

You can configure the console for bare metal installations in the following ways:

- Using **coreos-installer** manually on the command line.
- Using the **coreos-installer iso customize** or **coreos-installer pxe customize** subcommands with the **--dest-console** option to create a custom image that automates the process.



NOTE

For advanced customization, perform console configuration using the **coreos-installer iso** or **coreos-installer pxe** subcommands, and not kernel arguments.

15.4.11.3.5. Enabling the serial console for PXE and ISO installations

By default, the Red Hat Enterprise Linux CoreOS (RHCOS) serial console is disabled and all output is written to the graphical console. You can enable the serial console for an ISO installation and reconfigure the bootloader so that output is sent to both the serial console and the graphical console.

Procedure

1. Boot the ISO installer.
2. Run the **coreos-installer** command to install the system, adding the **--console** option once to specify the graphical console, and a second time to specify the serial console:

```
$ coreos-installer install \
  --console=tty0 \ 1
  --console=ttyS0,<options> \ 2
  --ignition-url=http://host/worker.ign /dev/disk/by-id/scsi-<serial_number>
```

- 1 The desired secondary console. In this case, the graphical console. Omitting this option will disable the graphical console.
- 2 The desired primary console. In this case the serial console. The **options** field defines the baud rate and other settings. A common value for this field is **11520n8**. If no options are provided, the default kernel value of **9600n8** is used. For more information on the format of this option, see [Linux kernel serial console](#) documentation.

3. Reboot into the installed system.



NOTE

A similar outcome can be obtained by using the **coreos-installer install --append-karg** option, and specifying the console with **console=**. However, this will only set the console for the kernel and not the bootloader.

To configure a PXE installation, make sure the **coreos.inst.install_dev** kernel command line option is omitted, and use the shell prompt to run **coreos-installer** manually using the above ISO installation procedure.

15.4.11.3.6. Customizing a live RHCOS ISO or PXE install

You can use the live ISO image or PXE environment to install RHCOS by injecting an Ignition config file directly into the image. This creates a customized image that you can use to provision your system.

For an ISO image, the mechanism to do this is the **coreos-installer iso customize** subcommand, which modifies the **.iso** file with your configuration. Similarly, the mechanism for a PXE environment is the **coreos-installer pxe customize** subcommand, which creates a new **initramfs** file that includes your customizations.

The **customize** subcommand is a general purpose tool that can embed other types of customizations as well. The following tasks are examples of some of the more common customizations:

- Inject custom CA certificates for when corporate security policy requires their use.
- Configure network settings without the need for kernel arguments.
- Embed arbitrary preinstall and post-install scripts or binaries.

15.4.11.3.7. Customizing a live RHCOS ISO image

You can customize a live RHCOS ISO image directly with the **coreos-installer iso customize** subcommand. When you boot the ISO image, the customizations are applied automatically.

You can use this feature to configure the ISO image to automatically install RHCOS.

Procedure

1. Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.
2. Retrieve the RHCOS ISO image from the [RHCOS image mirror](#) page and the Ignition config file, and then run the following command to inject the Ignition config directly into the ISO image:

```
$ coreos-installer iso customize rhcos-<version>-live.x86_64.iso \  
  --dest-ignition bootstrap.ign \ 1  
  --dest-device /dev/disk/by-id/scsi-<serial_number> 2
```

- 1 The Ignition config file that is generated from the **openshift-installer** installation program.
- 2 When you specify this option, the ISO image automatically runs an installation. Otherwise, the image remains configured for installation, but does not install automatically unless you specify the **coreos.inst.install_dev** kernel argument.

3. Optional: To remove the ISO image customizations and return the image to its pristine state, run:

```
$ coreos-installer iso reset rhcos-<version>-live.x86_64.iso
```

You can now re-customize the live ISO image or use it in its pristine state.

Applying your customizations affects every subsequent boot of RHCOS.

15.4.11.3.7.1. Modifying a live install ISO image to enable the serial console


On clusters installed with OpenShift Container Platform 4.12 and above, the serial console is disabled by default and all output is written to the graphical console. You can enable the serial console with the following procedure.

Procedure

1. Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.
2. Retrieve the RHCOS ISO image from the [RHCOS image mirror](#) page and run the following command to customize the ISO image to enable the serial console to receive output:

```
$ coreos-installer iso customize rhcos-<version>-live.x86_64.iso \
  --dest-ignition <path> \ 1
  --dest-console tty0 \ 2
  --dest-console ttyS0,<options> \ 3
  --dest-device /dev/disk/by-id/scsi-<serial_number> \ 4
```

- 1** The location of the Ignition config to install.
- 2** The desired secondary console. In this case, the graphical console. Omitting this option will disable the graphical console.
- 3** The desired primary console. In this case, the serial console. The **options** field defines the baud rate and other settings. A common value for this field is **115200n8**. If no options are provided, the default kernel value of **9600n8** is used. For more information on the format of this option, see the [Linux kernel serial console](#) documentation.
- 4** The specified disk to install to. If you omit this option, the ISO image automatically runs the installation program which will fail unless you also specify the **coreos.inst.install_dev** kernel argument.



NOTE

The **--dest-console** option affects the installed system and not the live ISO system. To modify the console for a live ISO system, use the **--live-karg-append** option and specify the console with **console=**.

Your customizations are applied and affect every subsequent boot of the ISO image.

3. Optional: To remove the ISO image customizations and return the image to its original state, run the following command:

```
$ coreos-installer iso reset rhcos-<version>-live.x86_64.iso
```

You can now recustomize the live ISO image or use it in its original state.

15.4.11.3.7.2. Modifying a live install ISO image to use a custom certificate authority

You can provide certificate authority (CA) certificates to Ignition with the **--ignition-ca** flag of the **customize** subcommand. You can use the CA certificates during both the installation boot and when provisioning the installed system.




NOTE

Custom CA certificates affect how Ignition fetches remote resources but they do not affect the certificates installed onto the system.

Procedure

1. Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.
2. Retrieve the RHCOS ISO image from the [RHCOS image mirror](#) page and run the following command to customize the ISO image for use with a custom CA:

```
$ coreos-installer iso customize rhcos-<version>-live.x86_64.iso --ignition-ca cert.pem
```



IMPORTANT

The **coreos.inst.ignition_url** kernel parameter does not work with the **--ignition-ca** flag. You must use the **--dest-ignition** flag to create a customized image for each cluster.

Applying your custom CA certificate affects every subsequent boot of RHCOS.

15.4.11.3.7.3. Modifying a live install ISO image with customized network settings

You can embed a NetworkManager keyfile into the live ISO image and pass it through to the installed system with the **--network-keyfile** flag of the **customize** subcommand.



WARNING

When creating a connection profile, you must use a **.nmconnection** filename extension in the filename of the connection profile. If you do not use a **.nmconnection** filename extension, the cluster will apply the connection profile to the live environment, but it will not apply the configuration when the cluster first boots up the nodes, resulting in a setup that does not work.

Procedure

1. Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.
2. Create a connection profile for a bonded interface. For example, create the **bond0.nmconnection** file in your local directory with the following content:

```
[connection]
id=bond0
type=bond
interface-name=bond0
```



```

multi-connect=1
permissions=

[ethernet]
mac-address-blacklist=

[bond]
miimon=100
mode=active-backup

[ipv4]
method=auto

[ipv6]
method=auto

[proxy]

```

3. Create a connection profile for a secondary interface to add to the bond. For example, create the **bond0-proxy-em1.nmconnection** file in your local directory with the following content:

```

[connection]
id=em1
type=ethernet
interface-name=em1
master=bond0
multi-connect=1
permissions=
slave-type=bond

[ethernet]
mac-address-blacklist=

```

4. Create a connection profile for a secondary interface to add to the bond. For example, create the **bond0-proxy-em2.nmconnection** file in your local directory with the following content:

```

[connection]
id=em2
type=ethernet
interface-name=em2
master=bond0
multi-connect=1
permissions=
slave-type=bond

[ethernet]
mac-address-blacklist=

```

5. Retrieve the RHCOS ISO image from the [RHCOS image mirror](#) page and run the following command to customize the ISO image with your configured networking:

```

$ coreos-installer iso customize rhcos-<version>-live.x86_64.iso \
  --network-keyfile bond0.nmconnection \
  --network-keyfile bond0-proxy-em1.nmconnection \
  --network-keyfile bond0-proxy-em2.nmconnection

```

Network settings are applied to the live system and are carried over to the destination system.

15.4.11.3.7.4. Customizing a live install ISO image for an iSCSI boot device

You can set the iSCSI target and initiator values for automatic mounting, booting and configuration using a customized version of the live RHCOS image.

Prerequisites

1. You have an iSCSI target you want to install RHCOS on.

Procedure

1. Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.
2. Retrieve the RHCOS ISO image from the [RHCOS image mirror](#) page and run the following command to customize the ISO image with the following information:

```
$ coreos-installer iso customize \
  --pre-install mount-iscsi.sh \ 1
  --post-install unmount-iscsi.sh \ 2
  --dest-device /dev/disk/by-path/<IP_address>:<port>-iscsi-<target_iqn>-lun-<lun> \ 3
  --dest-ignition config.ign \ 4
  --dest-karg-append rd.iscsi.initiator=<initiator_iqn> \ 5
  --dest-karg-append netroot=<target_iqn> \ 6
  -o custom.iso rhcos-<version>-live.x86_64.iso
```

- 1 The script that gets run before installation. It should contain the **iscsiadm** commands for mounting the iSCSI target and any commands enabling multipathing.
- 2 The script that gets run after installation. It should contain the command **iscsiadm --mode node --logout=all**.
- 3 The location of the destination system. You must provide the IP address of the target portal, the associated port number, the target iSCSI node in IQN format, and the iSCSI logical unit number (LUN).
- 4 The Ignition configuration for the destination system.
- 5 The iSCSI initiator, or client, name in IQN format. The initiator forms a session to connect to the iSCSI target.
- 6 The the iSCSI target, or server, name in IQN format.

For more information about the iSCSI options supported by **dracut**, see the [dracut.cmdline manual page](#).

15.4.11.3.7.5. Customizing a live install ISO image for an iSCSI boot device with iBFT

You can set the iSCSI target and initiator values for automatic mounting, booting and configuration using a customized version of the live RHCOS image.

Prerequisites

1. You have an iSCSI target you want to install RHCOS on.
2. Optional: you have multipathed your iSCSI target.

Procedure

1. Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.
2. Retrieve the RHCOS ISO image from the [RHCOS image mirror](#) page and run the following command to customize the ISO image with the following information:

```
$ coreos-installer iso customize \
  --pre-install mount-iscsi.sh \ 1
  --post-install unmount-iscsi.sh \ 2
  --dest-device /dev/mapper/mpatha \ 3
  --dest-ignition config.ign \ 4
  --dest-karg-append rd.iscsi.firmware=1 \ 5
  --dest-karg-append rd.multipath=default \ 6
  -o custom.iso rhcos-<version>-live.x86_64.iso
```

- 1 The script that gets run before installation. It should contain the **iscsiadm** commands for mounting the iSCSI target and any commands enabling multipathing.
- 2 The script that gets run after installation. It should contain the command **iscsiadm --mode node --logout=all**.
- 3 The path to the device. If you are using multipath, the multipath device, **/dev/mapper/mpatha**, If there are multiple multipath devices connected, or to be explicit, you can use the World Wide Name (WWN) symlink available in **/dev/disk/by-path**.
- 4 The Ignition configuration for the destination system.
- 5 The iSCSI parameter is read from the BIOS firmware.
- 6 Optional: include this parameter if you are enabling multipathing.

For more information about the iSCSI options supported by **dracut**, see the [dracut.cmdline manual page](#).

15.4.11.3.8. Customizing a live RHCOS PXE environment

You can customize a live RHCOS PXE environment directly with the **coreos-installer pxe customize** subcommand. When you boot the PXE environment, the customizations are applied automatically.

You can use this feature to configure the PXE environment to automatically install RHCOS.

Procedure

1. Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.
2. Retrieve the RHCOS **kernel**, **initramfs** and **rootfs** files from the [RHCOS image mirror](#) page and the Ignition config file, and then run the following command to create a new **initramfs** file that contains the customizations from your Ignition config:

```
$ coreos-installer pxe customize rhcos-<version>-live-initramfs.x86_64.img \
  --dest-ignition bootstrap.ign \ 1
  --dest-device /dev/disk/by-id/scsi-<serial_number> \ 2
  -o rhcos-<version>-custom-initramfs.x86_64.img 3
```

- 1 The Ignition config file that is generated from **openshift-installer**.
- 2 When you specify this option, the PXE environment automatically runs an install. Otherwise, the image remains configured for installing, but does not do so automatically unless you specify the **coreos.inst.install_dev** kernel argument.
- 3 Use the customized **initramfs** file in your PXE configuration. Add the **ignition.firstboot** and **ignition.platform.id=metal** kernel arguments if they are not already present.

Applying your customizations affects every subsequent boot of RHCOS.

15.4.11.3.8.1. Modifying a live install PXE environment to enable the serial console

On clusters installed with OpenShift Container Platform 4.12 and above, the serial console is disabled by default and all output is written to the graphical console. You can enable the serial console with the following procedure.

Procedure

1. Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.
2. Retrieve the RHCOS **kernel**, **initramfs** and **rootfs** files from the [RHCOS image mirror](#) page and the Ignition config file, and then run the following command to create a new customized **initramfs** file that enables the serial console to receive output:

```
$ coreos-installer pxe customize rhcos-<version>-live-initramfs.x86_64.img \
  --dest-ignition <path> \ 1
  --dest-console tty0 \ 2
  --dest-console ttyS0,<options> \ 3
  --dest-device /dev/disk/by-id/scsi-<serial_number> \ 4
  -o rhcos-<version>-custom-initramfs.x86_64.img 5
```

- 1 The location of the Ignition config to install.
- 2 The desired secondary console. In this case, the graphical console. Omitting this option will disable the graphical console.
- 3 The desired primary console. In this case, the serial console. The **options** field defines the baud rate and other settings. A common value for this field is **115200n8**. If no options are provided, the default kernel value of **9600n8** is used. For more information on the format of this option, see the [Linux kernel serial console](#) documentation.
- 4 The specified disk to install to. If you omit this option, the PXE environment automatically runs the installer which will fail unless you also specify the **coreos.inst.install_dev** kernel argument.
- 5 Use the customized **initramfs** file in your PXE configuration. Add the **ignition.firstboot** and **ignition.platform.id=metal** kernel arguments if they are not already present.

Your customizations are applied and affect every subsequent boot of the PXE environment.

15.4.11.3.8.2. Modifying a live install PXE environment to use a custom certificate authority

You can provide certificate authority (CA) certificates to Ignition with the **--ignition-ca** flag of the **customize** subcommand. You can use the CA certificates during both the installation boot and when provisioning the installed system.



NOTE

Custom CA certificates affect how Ignition fetches remote resources but they do not affect the certificates installed onto the system.

Procedure

1. Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.
2. Retrieve the RHCOS **kernel**, **initramfs** and **rootfs** files from the [RHCOS image mirror](#) page and run the following command to create a new customized **initramfs** file for use with a custom CA:

```
$ coreos-installer pxe customize rhcos-<version>-live-initramfs.x86_64.img \
--ignition-ca cert.pem \
-o rhcos-<version>-custom-initramfs.x86_64.img
```

3. Use the customized **initramfs** file in your PXE configuration. Add the **ignition.firstboot** and **ignition.platform.id=metal** kernel arguments if they are not already present.



IMPORTANT

The **coreos.inst.ignition_url** kernel parameter does not work with the **--ignition-ca** flag. You must use the **--dest-ignition** flag to create a customized image for each cluster.

Applying your custom CA certificate affects every subsequent boot of RHCOS.

15.4.11.3.8.3. Modifying a live install PXE environment with customized network settings

You can embed a NetworkManager keyfile into the live PXE environment and pass it through to the installed system with the **--network-keyfile** flag of the **customize** subcommand.



WARNING

When creating a connection profile, you must use a **.nmconnection** filename extension in the filename of the connection profile. If you do not use a **.nmconnection** filename extension, the cluster will apply the connection profile to the live environment, but it will not apply the configuration when the cluster first boots up the nodes, resulting in a setup that does not work.

Procedure

1. Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.
2. Create a connection profile for a bonded interface. For example, create the **bond0.nmconnection** file in your local directory with the following content:

```
[connection]
id=bond0
type=bond
interface-name=bond0
multi-connect=1
permissions=

[ethernet]
mac-address-blacklist=

[bond]
miimon=100
mode=active-backup

[ipv4]
method=auto

[ipv6]
method=auto

[proxy]
```

3. Create a connection profile for a secondary interface to add to the bond. For example, create the **bond0-proxy-em1.nmconnection** file in your local directory with the following content:

```
[connection]
id=em1
type=ethernet
interface-name=em1
master=bond0
multi-connect=1
permissions=
slave-type=bond

[ethernet]
mac-address-blacklist=
```

4. Create a connection profile for a secondary interface to add to the bond. For example, create the **bond0-proxy-em2.nmconnection** file in your local directory with the following content:

```
[connection]
id=em2
type=ethernet
interface-name=em2
master=bond0
multi-connect=1
permissions=
slave-type=bond
```

```
[ethernet]
mac-address-blacklist=
```

- Retrieve the RHCOS **kernel**, **initramfs** and **rootfs** files from the [RHCOS image mirror](#) page and run the following command to create a new customized **initramfs** file that contains your configured networking:

```
$ coreos-installer pxe customize rhcos-<version>-live-initramfs.x86_64.img \
  --network-keyfile bond0.nmconnection \
  --network-keyfile bond0-proxy-em1.nmconnection \
  --network-keyfile bond0-proxy-em2.nmconnection \
  -o rhcos-<version>-custom-initramfs.x86_64.img
```

- Use the customized **initramfs** file in your PXE configuration. Add the **ignition.firstboot** and **ignition.platform.id=metal** kernel arguments if they are not already present. Network settings are applied to the live system and are carried over to the destination system.

15.4.11.3.8.4. Customizing a live install PXE environment for an iSCSI boot device

You can set the iSCSI target and initiator values for automatic mounting, booting and configuration using a customized version of the live RHCOS image.

Prerequisites

- You have an iSCSI target you want to install RHCOS on.

Procedure

- Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.
- Retrieve the RHCOS **kernel**, **initramfs** and **rootfs** files from the [RHCOS image mirror](#) page and run the following command to create a new customized **initramfs** file with the following information:

```
$ coreos-installer pxe customize \
  --pre-install mount-iscsi.sh \ 1
  --post-install unmount-iscsi.sh \ 2
  --dest-device /dev/disk/by-path/<IP_address>:<port>-iscsi-<target_iqn>-lun-<lun> \ 3
  --dest-ignition config.ign \ 4
  --dest-karg-append rd.iscsi.initiator=<initiator_iqn> \ 5
  --dest-karg-append netroot=<target_iqn> \ 6
  -o custom.img rhcos-<version>-live-initramfs.x86_64.img
```

- The script that gets run before installation. It should contain the **iscsiadm** commands for mounting the iSCSI target and any commands enabling multipathing.
- The script that gets run after installation. It should contain the command **iscsiadm --mode node --logout=all**.
- The location of the destination system. You must provide the IP address of the target portal, the associated port number, the target iSCSI node in IQN format, and the iSCSI logical unit number (LUN).

- 4 The Ignition configuration for the destination system.
- 5 The iSCSI initiator, or client, name in IQN format. The initiator forms a session to connect to the iSCSI target.
- 6 The the iSCSI target, or server, name in IQN format.

For more information about the iSCSI options supported by **dracut**, see the [dracut.cmdline manual page](#).

15.4.11.3.8.5. Customizing a live install PXE environment for an iSCSI boot device with iBFT

You can set the iSCSI target and initiator values for automatic mounting, booting and configuration using a customized version of the live RHCOS image.

Prerequisites

1. You have an iSCSI target you want to install RHCOS on.
2. Optional: you have multipathed your iSCSI target.

Procedure

1. Download the **coreos-installer** binary from the [coreos-installer image mirror](#) page.
2. Retrieve the RHCOS **kernel**, **initramfs** and **rootfs** files from the [RHCOS image mirror](#) page and run the following command to create a new customized **initramfs** file with the following information:

```
$ coreos-installer pxe customize \
  --pre-install mount-iscsi.sh \ 1
  --post-install unmount-iscsi.sh \ 2
  --dest-device /dev/mapper/mpatha \ 3
  --dest-ignition config.ign \ 4
  --dest-karg-append rd.iscsi.firmware=1 \ 5
  --dest-karg-append rd.multipath=default \ 6
  -o custom.img rhcos-<version>-live-initramfs.x86_64.img
```

- 1 The script that gets run before installation. It should contain the **iscsiadm** commands for mounting the iSCSI target.
- 2 The script that gets run after installation. It should contain the command **iscsiadm --mode node --logout=all**.
- 3 The path to the device. If you are using multipath, the multipath device, **/dev/mapper/mpatha**, If there are multiple multipath devices connected, or to be explicit, you can use the World Wide Name (WWN) symlink available in **/dev/disk/by-path**.
- 4 The Ignition configuration for the destination system.
- 5 The iSCSI parameter is read from the BIOS firmware.
- 6 Optional: include this parameter if you are enabling multipathing.

For more information about the iSCSI options supported by **dracut**, see the [dracut.cmdline manual page](#).

15.4.11.3.9. Advanced RHCOS installation reference

This section illustrates the networking configuration and other advanced options that allow you to modify the Red Hat Enterprise Linux CoreOS (RHCOS) manual installation process. The following tables describe the kernel arguments and command-line options you can use with the RHCOS live installer and the **coreos-installer** command.

15.4.11.3.9.1. Networking and bonding options for ISO installations

If you install RHCOS from an ISO image, you can add kernel arguments manually when you boot the image to configure networking for a node. If no networking arguments are specified, DHCP is activated in the initramfs when RHCOS detects that networking is required to fetch the Ignition config file.



IMPORTANT

When adding networking arguments manually, you must also add the **rd.neednet=1** kernel argument to bring the network up in the initramfs.

The following information provides examples for configuring networking and bonding on your RHCOS nodes for ISO installations. The examples describe how to use the **ip=**, **nameserver=**, and **bond=** kernel arguments.



NOTE

Ordering is important when adding the kernel arguments: **ip=**, **nameserver=**, and then **bond=**.

The networking options are passed to the **dracut** tool during system boot. For more information about the networking options supported by **dracut**, see the [dracut.cmdline manual page](#).

The following examples are the networking options for ISO installation.

Configuring DHCP or static IP addresses

To configure an IP address, either use DHCP (**ip=dhcp**) or set an individual static IP address (**ip=<host_ip>**). If setting a static IP, you must then identify the DNS server IP address (**nameserver=<dns_ip>**) on each node. The following example sets:

- The node's IP address to **10.10.10.2**
- The gateway address to **10.10.10.254**
- The netmask to **255.255.255.0**
- The hostname to **core0.example.com**
- The DNS server address to **4.4.4.41**
- The auto-configuration value to **none**. No auto-configuration is required when IP networking is configured statically.

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none
nameserver=4.4.4.41
```



NOTE

When you use DHCP to configure IP addressing for the RHCOS machines, the machines also obtain the DNS server information through DHCP. For DHCP-based deployments, you can define the DNS server address that is used by the RHCOS nodes through your DHCP server configuration.

Configuring an IP address without a static hostname

You can configure an IP address without assigning a static hostname. If a static hostname is not set by the user, it will be picked up and automatically set by a reverse DNS lookup. To configure an IP address without a static hostname refer to the following example:

- The node's IP address to **10.10.10.2**
- The gateway address to **10.10.10.254**
- The netmask to **255.255.255.0**
- The DNS server address to **4.4.4.41**
- The auto-configuration value to **none**. No auto-configuration is required when IP networking is configured statically.

```
ip=10.10.10.2::10.10.10.254:255.255.255.0::enp1s0:none
nameserver=4.4.4.41
```

Specifying multiple network interfaces

You can specify multiple network interfaces by setting multiple **ip=** entries.

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none
ip=10.10.10.3::10.10.10.254:255.255.255.0:core0.example.com:enp2s0:none
```

Configuring default gateway and route

Optional: You can configure routes to additional networks by setting an **rd.route=** value.



NOTE

When you configure one or multiple networks, one default gateway is required. If the additional network gateway is different from the primary network gateway, the default gateway must be the primary network gateway.

- Run the following command to configure the default gateway:

```
ip=::10.10.10.254:::
```

- Enter the following command to configure the route for the additional network:

```
rd.route=20.20.20.0/24:20.20.20.254:enp2s0
```

Disabling DHCP on a single interface

You can disable DHCP on a single interface, such as when there are two or more network interfaces and only one interface is being used. In the example, the **enp1s0** interface has a static networking configuration and DHCP is disabled for **enp2s0**, which is not used:

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none
ip=::::core0.example.com:enp2s0:none
```

Combining DHCP and static IP configurations

You can combine DHCP and static IP configurations on systems with multiple network interfaces, for example:

```
ip=enp1s0:dhcp
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp2s0:none
```

Configuring VLANs on individual interfaces

Optional: You can configure VLANs on individual interfaces by using the **vlan=** parameter.

- To configure a VLAN on a network interface and use a static IP address, run the following command:

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp2s0.100:none
vlan=enp2s0.100:enp2s0
```

- To configure a VLAN on a network interface and to use DHCP, run the following command:

```
ip=enp2s0.100:dhcp
vlan=enp2s0.100:enp2s0
```

Providing multiple DNS servers

You can provide multiple DNS servers by adding a **nameserver=** entry for each server, for example:

```
nameserver=1.1.1.1
nameserver=8.8.8.8
```

Bonding multiple network interfaces to a single interface

Optional: You can bond multiple network interfaces to a single interface by using the **bond=** option.

Refer to the following examples:


- The syntax for configuring a bonded interface is: **bond=<name>[:<network_interfaces>][:options]**
<name> is the bonding device name (**bond0**), **<network_interfaces>** represents a comma-separated list of physical (ethernet) interfaces (**em1,em2**), and *options* is a comma-separated list of bonding options. Enter **modinfo bonding** to see available options.
- When you create a bonded interface using **bond=**, you must specify how the IP address is assigned and other information for the bonded interface.
 - To configure the bonded interface to use DHCP, set the bond's IP address to **dhcp**. For example:

```
bond=bond0:em1,em2:mode=active-backup
ip=bond0:dhcp
```

- To configure the bonded interface to use a static IP address, enter the specific IP address you want and related information. For example:

```
bond=bond0:em1,em2:mode=active-backup
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:bond0:none
```

Bonding multiple SR-IOV network interfaces to a dual port NIC interface



IMPORTANT

Support for Day 1 operations associated with enabling NIC partitioning for SR-IOV devices is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

Optional: You can bond multiple SR-IOV network interfaces to a dual port NIC interface by using the **bond=** option.

On each node, you must perform the following tasks:

1. Create the SR-IOV virtual functions (VFs) following the guidance in [Managing SR-IOV devices](#). Follow the procedure in the "Attaching SR-IOV networking devices to virtual machines" section.
2. Create the bond, attach the desired VFs to the bond and set the bond link state up following the guidance in [Configuring network bonding](#). Follow any of the described procedures to create the bond.

The following examples illustrate the syntax you must use:

- The syntax for configuring a bonded interface is **bond=<name>[:<network_interfaces>][:options]**. **<name>** is the bonding device name (**bond0**), **<network_interfaces>** represents the virtual functions (VFs) by their known name in the kernel and shown in the output of the **ip link** command (**eno1f0**, **eno2f0**), and **options** is a comma-separated list of bonding options. Enter **modinfo bonding** to see available options.
- When you create a bonded interface using **bond=**, you must specify how the IP address is assigned and other information for the bonded interface.
 - To configure the bonded interface to use DHCP, set the bond's IP address to **dhcp**. For example:

```
bond=bond0:eno1f0,eno2f0:mode=active-backup
ip=bond0:dhcp
```

- To configure the bonded interface to use a static IP address, enter the specific IP address you want and related information. For example:

```
bond=bond0:eno1f0,eno2f0:mode=active-backup
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:bond0:none
```

Using network teaming

Optional: You can use a network teaming as an alternative to bonding by using the **team=** parameter:

- The syntax for configuring a team interface is: **team=name[:network_interfaces]**
name is the team device name (**team0**) and *network_interfaces* represents a comma-separated list of physical (ethernet) interfaces (**em1, em2**).



NOTE

Teaming is planned to be deprecated when RHCOS switches to an upcoming version of RHEL. For more information, see this [Red Hat Knowledgebase Article](#).

Use the following example to configure a network team:

```
team=team0:em1,em2
ip=team0:dhcp
```


15.4.11.3.9.2. coreos-installer options for ISO and PXE installations

You can install RHCOS by running **coreos-installer install <options> <device>** at the command prompt, after booting into the RHCOS live environment from an ISO image.

The following table shows the subcommands, options, and arguments you can pass to the **coreos-installer** command.


Table 15.40. coreos-installer subcommands, command-line options, and arguments

coreos-installer install subcommand	
<i>Subcommand</i>	<i>Description</i>
\$ coreos-installer install <options> <device>	Embed an Ignition config in an ISO image.
coreos-installer install subcommand options	
<i>Option</i>	<i>Description</i>
-u, --image-url <url>	Specify the image URL manually.
-f, --image-file <path>	Specify a local image file manually. Used for debugging.
-i, --ignition-file <path>	Embed an Ignition config from a file.
-l, --ignition-url <URL>	Embed an Ignition config from a URL.
--ignition-hash <digest>	Digest type-value of the Ignition config.
-p, --platform <name>	Override the Ignition platform ID for the installed system.

--console <spec>	Set the kernel and bootloader console for the installed system. For more information about the format of <spec> , see the Linux kernel serial console documentation.
--append-karg <arg>...	Append a default kernel argument to the installed system.
--delete-karg <arg>...	Delete a default kernel argument from the installed system.
-n, --copy-network	Copy the network configuration from the install environment.  IMPORTANT The --copy-network option only copies networking configuration found under /etc/NetworkManager/system-connections . In particular, it does not copy the system hostname.
--network-dir <path>	For use with -n . Default is /etc/NetworkManager/system-connections/ .
--save-partlabel <lx>..	Save partitions with this label glob.
--save-partindex <id>...	Save partitions with this number or range.
--insecure	Skip RHCOS image signature verification.
--insecure-ignition	Allow Ignition URL without HTTPS or hash.
--architecture <name>	Target CPU architecture. Valid values are x86_64 and aarch64 .
--preserve-on-error	Do not clear partition table on error.
-h, --help	Print help information.
coreos-installer install subcommand argument	
<i>Argument</i>	<i>Description</i>
<device>	The destination device.
coreos-installer ISO subcommands	

<i>Subcommand</i>	<i>Description</i>
\$ coreos-installer iso customize <options> <ISO_image>	Customize a RHCOS live ISO image.
coreos-installer iso reset <options> <ISO_image>	Restore a RHCOS live ISO image to default settings.
coreos-installer iso ignition remove <options> <ISO_image>	Remove the embedded Ignition config from an ISO image.
coreos-installer ISO customize subcommand options	
<i>Option</i>	<i>Description</i>
--dest-ignition <path>	Merge the specified Ignition config file into a new configuration fragment for the destination system.
--dest-console <spec>	Specify the kernel and bootloader console for the destination system.
--dest-device <path>	Install and overwrite the specified destination device.
--dest-karg-append <arg>	Add a kernel argument to each boot of the destination system.
--dest-karg-delete <arg>	Delete a kernel argument from each boot of the destination system.
--network-keyfile <path>	Configure networking by using the specified NetworkManager keyfile for live and destination systems.
--ignition-ca <path>	Specify an additional TLS certificate authority to be trusted by Ignition.
--pre-install <path>	Run the specified script before installation.
--post-install <path>	Run the specified script after installation.
--installer-config <path>	Apply the specified installer configuration file.
--live-ignition <path>	Merge the specified Ignition config file into a new configuration fragment for the live environment.
--live-karg-append <arg>	Add a kernel argument to each boot of the live environment.

--live-karg-delete <arg>	Delete a kernel argument from each boot of the live environment.
--live-karg-replace <k=o=n>	Replace a kernel argument in each boot of the live environment, in the form key=old=new .
-f, --force	Overwrite an existing Ignition config.
-o, --output <path>	Write the ISO to a new output file.
-h, --help	Print help information.
coreos-installer PXE subcommands	
<i>Subcommand</i>	<i>Description</i>
Note that not all of these options are accepted by all subcommands.	
coreos-installer pxe customize <options> <path>	Customize a RHCOS live PXE boot config.
coreos-installer pxe ignition wrap <options>	Wrap an Ignition config in an image.
coreos-installer pxe ignition unwrap <options> <image_name>	Show the wrapped Ignition config in an image.
coreos-installer PXE customize subcommand options	
<i>Option</i>	<i>Description</i>
Note that not all of these options are accepted by all subcommands.	
--dest-ignition <path>	Merge the specified Ignition config file into a new configuration fragment for the destination system.
--dest-console <spec>	Specify the kernel and bootloader console for the destination system.
--dest-device <path>	Install and overwrite the specified destination device.
--network-keyfile <path>	Configure networking by using the specified NetworkManager keyfile for live and destination systems.
--ignition-ca <path>	Specify an additional TLS certificate authority to be trusted by Ignition.
--pre-install <path>	Run the specified script before installation.

post-install <path>	Run the specified script after installation.
--installer-config <path>	Apply the specified installer configuration file.
--live-ignition <path>	Merge the specified Ignition config file into a new configuration fragment for the live environment.
-o, --output <path>	Write the initramfs to a new output file.
	 <p>NOTE</p> <p>This option is required for PXE environments.</p>
-h, --help	Print help information.

15.4.11.3.9.3. coreos.inst boot options for ISO or PXE installations

You can automatically invoke **coreos-installer** options at boot time by passing **coreos.inst** boot arguments to the RHCOS live installer. These are provided in addition to the standard boot arguments.

- For ISO installations, the **coreos.inst** options can be added by interrupting the automatic boot at the bootloader menu. You can interrupt the automatic boot by pressing **TAB** while the **RHEL CoreOS (Live)** menu option is highlighted.
- For PXE or iPXE installations, the **coreos.inst** options must be added to the **APPEND** line before the RHCOS live installer is booted.

The following table shows the RHCOS live installer **coreos.inst** boot options for ISO and PXE installations.

Table 15.41. **coreos.inst** boot options

Argument	Description
coreos.inst.install_dev	Required. The block device on the system to install to. It is recommended to use the full path, such as /dev/sda , although sda is allowed.
coreos.inst.ignition_url	Optional: The URL of the Ignition config to embed into the installed system. If no URL is specified, no Ignition config is embedded. Only HTTP and HTTPS protocols are supported.
coreos.inst.save_partlabel	Optional: Comma-separated labels of partitions to preserve during the install. Glob-style wildcards are permitted. The specified partitions do not need to exist.

Argument	Description
coreos.inst.save_partindex	Optional: Comma-separated indexes of partitions to preserve during the install. Ranges m-n are permitted, and either m or n can be omitted. The specified partitions do not need to exist.
coreos.inst.insecure	Optional: Permits the OS image that is specified by coreos.inst.image_url to be unsigned.
coreos.inst.image_url	<p>Optional: Download and install the specified RHCOS image.</p> <ul style="list-style-type: none"> ● This argument should not be used in production environments and is intended for debugging purposes only. ● While this argument can be used to install a version of RHCOS that does not match the live media, it is recommended that you instead use the media that matches the version you want to install. ● If you are using coreos.inst.image_url, you must also use coreos.inst.insecure. This is because the bare-metal media are not GPG-signed for OpenShift Container Platform. ● Only HTTP and HTTPS protocols are supported.
coreos.inst.skip_reboot	Optional: The system will not reboot after installing. After the install finishes, you will receive a prompt that allows you to inspect what is happening during installation. This argument should not be used in production environments and is intended for debugging purposes only.
coreos.inst.platform_id	Optional: The Ignition platform ID of the platform the RHCOS image is being installed on. Default is metal . This option determines whether or not to request an Ignition config from the cloud provider, such as VMware. For example: coreos.inst.platform_id=vmware .
ignition.config.url	Optional: The URL of the Ignition config for the live boot. For example, this can be used to customize how coreos-installer is invoked, or to run code before or after the installation. This is different from coreos.inst.ignition_url , which is the Ignition config for the installed system.

15.4.11.4. Enabling multipathing with kernel arguments on RHCOS

RHCOS supports multipathing on the primary disk, allowing stronger resilience to hardware failure to achieve higher host availability.

You can enable multipathing at installation time for nodes that were provisioned in OpenShift Container Platform 4.8 or later. While postinstallation support is available by activating multipathing via the machine config, enabling multipathing during installation is recommended.

In setups where any I/O to non-optimized paths results in I/O system errors, you must enable multipathing at installation time.



IMPORTANT

On IBM Z® and IBM® LinuxONE, you can enable multipathing only if you configured your cluster for it during installation. For more information, see "Installing RHCOS and starting the OpenShift Container Platform bootstrap process" in *Installing a cluster with z/VM on IBM Z® and IBM® LinuxONE*.

The following procedure enables multipath at installation time and appends kernel arguments to the **coreos-installer install** command so that the installed system itself will use multipath beginning from the first boot.



NOTE

OpenShift Container Platform does not support enabling multipathing as a day-2 activity on nodes that have been upgraded from 4.6 or earlier.

Procedure

1. To enable multipath and start the **multipathd** daemon, run the following command on the installation host:

```
$ mpathconf --enable && systemctl start multipathd.service
```

- Optional: If booting the PXE or ISO, you can instead enable multipath by adding **rd.multipath=default** from the kernel command line.

2. Append the kernel arguments by invoking the **coreos-installer** program:

- If there is only one multipath device connected to the machine, it should be available at path **/dev/mapper/mpatha**. For example:

```
$ coreos-installer install /dev/mapper/mpatha \ 1
--append-karg rd.multipath=default \
--append-karg root=/dev/disk/by-label/dm-mpath-root \
--append-karg rw
```

- 1 Indicates the path of the single multipathed device.

- If there are multiple multipath devices connected to the machine, or to be more explicit, instead of using **/dev/mapper/mpatha**, it is recommended to use the World Wide Name (WWN) symlink available in **/dev/disk/by-id**. For example:

```
$ coreos-installer install /dev/disk/by-id/wwn-<wwn_ID> \ 1
--append-karg rd.multipath=default \
--append-karg root=/dev/disk/by-label/dm-mpath-root \
--append-karg rw
```

- 1** Indicates the WWN ID of the target multipathed device. For example, **0xx194e957fcedb4841**.

This symlink can also be used as the **coreos.inst.install_dev** kernel argument when using special **coreos.inst.*** arguments to direct the live installer. For more information, see "Installing RHCOS and starting the OpenShift Container Platform bootstrap process".

3. Check that the kernel arguments worked by going to one of the worker nodes and listing the kernel command line arguments (in **/proc/cmdline** on the host):

```
$ oc debug node/ip-10-0-141-105.ec2.internal
```

Example output

```
Starting pod/ip-10-0-141-105ec2internal-debug ...
To use host binaries, run `chroot /host`

sh-4.2# cat /host/proc/cmdline
...
rd.multipath=default root=/dev/disk/by-label/dm-mpath-root
...

sh-4.2# exit
```

You should see the added kernel arguments.

15.4.11.5. Installing RHCOS manually on an iSCSI boot device

You can manually install RHCOS on an iSCSI target.

Prerequisites

1. You are in the RHCOS live environment.
2. You have an iSCSI target that you want to install RHCOS on.

Procedure

1. Mount the iSCSI target from the live environment by running the following command:

```
$ iscsiadm \
--mode discovery \
--type sendtargets
--portal <IP_address> \ 1
--login
```

- 1** The IP address of the target portal.

2. Install RHCOS onto the iSCSI target by running the following command and using the necessary kernel arguments, for example:

```
$ coreos-installer install \
  /dev/disk/by-path/ip-<IP_address>:<port>-iscsi-<target_iqn>-lun-<lun> \ 1
  --append-karg rd.iscsi.initiator=<initiator_iqn> \ 2
  --append.karg netroot=<target_iqn> \ 3
  --console ttyS0,115200n8
  --ignition-file <path_to_file>
```

- 1 The location you are installing to. You must provide the IP address of the target portal, the associated port number, the target iSCSI node in IQN format, and the iSCSI logical unit number (LUN).
- 2 The iSCSI initiator, or client, name in IQN format. The initiator forms a session to connect to the iSCSI target.
- 3 The the iSCSI target, or server, name in IQN format.

For more information about the iSCSI options supported by **dracut**, see the [dracut.cmdline manual page](#).

3. Unmount the iSCSI disk with the following command:

```
$ iscsiadm --mode node --logoutall=all
```

This procedure can also be performed using the **coreos-installer iso customize** or **coreos-installer pxe customize** subcommands.

15.4.11.6. Installing RHCOS on an iSCSI boot device using iBFT

On a completely diskless machine, the iSCSI target and initiator values can be passed through iBFT. iSCSI multipathing is also supported.

Prerequisites

1. You are in the RHCOS live environment.
2. You have an iSCSI target you want to install RHCOS on.
3. Optional: you have multipathed your iSCSI target.

Procedure

1. Mount the iSCSI target from the live environment by running the following command:

```
$ iscsiadm \
  --mode discovery \
  --type sendtargets \
  --portal <IP_address> \ 1
  --login
```

- 1 The IP address of the target portal.

- Optional: enable multipathing and start the daemon with the following command:

```
$ mpathconf --enable && systemctl start multipathd.service
```

- Install RHCOS onto the iSCSI target by running the following command and using the necessary kernel arguments, for example:

```
$ coreos-installer install \  
  /dev/mapper/mpatha \ 1 \  
  --append-karg rd.iscsi.firmware=1 \ 2 \  
  --append-karg rd.multipath=default \ 3 \  
  --console ttyS0 \  
  --ignition-file <path_to_file>
```

1 The path of a single multipathed device. If there are multiple multipath devices connected, or to be explicit, you can use the World Wide Name (WWN) symlink available in **/dev/disk/by-path**.

2 The iSCSI parameter is read from the BIOS firmware.

3 Optional: include this parameter if you are enabling multipathing.

For more information about the iSCSI options supported by **dracut**, see the [dracut.cmdline manual page](#).

- Unmount the iSCSI disk:

```
$ iscsiadm --mode node --logout=all
```

This procedure can also be performed using the **coreos-installer iso customize** or **coreos-installer pxe customize** subcommands.

15.4.12. Waiting for the bootstrap process to complete

The OpenShift Container Platform bootstrap process begins after the cluster nodes first boot into the persistent RHCOS environment that has been installed to disk. The configuration information provided through the Ignition config files is used to initialize the bootstrap process and install OpenShift Container Platform on the machines. You must wait for the bootstrap process to complete.

Prerequisites

- You have created the Ignition config files for your cluster.
- You have configured suitable network, DNS and load balancing infrastructure.
- You have obtained the installation program and generated the Ignition config files for your cluster.
- You installed RHCOS on your cluster machines and provided the Ignition config files that the OpenShift Container Platform installation program generated.

Procedure

1. Monitor the bootstrap process:

```
$ ./openshift-install --dir <installation_directory> wait-for bootstrap-complete \ 1
--log-level=info 2
```

1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.


2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Example output

```
INFO Waiting up to 30m0s for the Kubernetes API at https://api.test.example.com:6443...
INFO API v1.29.4 up
INFO Waiting up to 30m0s for bootstrapping to complete...
INFO It is now safe to remove the bootstrap resources
```

The command succeeds when the Kubernetes API server signals that it has been bootstrapped on the control plane machines.

2. After the bootstrap process is complete, remove the bootstrap machine from the load balancer.



IMPORTANT

You must remove the bootstrap machine from the load balancer at this point. You can also remove or reformat the bootstrap machine itself.

Additional resources

- See [Monitoring installation progress](#) for more information about monitoring the installation logs and retrieving diagnostic data if installation issues arise.

15.4.13. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1 For `<installation_directory>`, specify the path to the directory that you stored the installation files in.

2. Verify you can run `oc` commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

15.4.14. Approving the certificate signing requests for your machines

When you add machines to a cluster, two pending certificate signing requests (CSRs) are generated for each machine that you added. You must confirm that these CSRs are approved or, if necessary, approve them yourself. The client requests must be approved first, followed by the server requests.

Prerequisites

- You added machines to your cluster.

Procedure

1. Confirm that the cluster recognizes the machines:

```
$ oc get nodes
```

Example output

```
NAME      STATUS   ROLES    AGE   VERSION
master-0  Ready   master   63m   v1.29.4
master-1  Ready   master   63m   v1.29.4
master-2  Ready   master   64m   v1.29.4
```

The output lists all of the machines that you created.



NOTE

The preceding output might not include the compute nodes, also known as worker nodes, until some CSRs are approved.

2. Review the pending CSRs and ensure that you see the client requests with the **Pending** or **Approved** status for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

```
NAME      AGE   REQUESTOR                                     CONDITION
csr-8b2br  15m   system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper  Pending
```



```
csr-8vnps 15m system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper Pending
...
```

In this example, two machines are joining the cluster. You might see more approved CSRs in the list.

- If the CSRs were not approved, after all of the pending CSRs for the machines you added are in **Pending** status, approve the CSRs for your cluster machines:



NOTE

Because the CSRs rotate automatically, approve your CSRs within an hour of adding the machines to the cluster. If you do not approve them within an hour, the certificates will rotate, and more than two certificates will be present for each node. You must approve all of these certificates. After the client CSR is approved, the Kubelet creates a secondary CSR for the serving certificate, which requires manual approval. Then, subsequent serving certificate renewal requests are automatically approved by the **machine-approver** if the Kubelet requests a new certificate with identical parameters.



NOTE

For clusters running on platforms that are not machine API enabled, such as bare metal and other user-provisioned infrastructure, you must implement a method of automatically approving the kubelet serving certificate requests (CSRs). If a request is not approved, then the **oc exec**, **oc rsh**, and **oc logs** commands cannot succeed, because a serving certificate is required when the API server connects to the kubelet. Any operation that contacts the Kubelet endpoint requires this certificate approval to be in place. The method must watch for new CSRs, confirm that the CSR was submitted by the **node-bootstrapper** service account in the **system:node** or **system:admin** groups, and confirm the identity of the node.

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{"\n"}
{{end}}{{end}}' | xargs --no-run-if-empty oc adm certificate approve
```



NOTE

Some Operators might not become available until some CSRs are approved.

- Now that your client requests are approved, you must review the server requests for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

```
NAME      AGE   REQUESTOR                                     CONDITION
csr-bfd72 5m26s system:node:ip-10-0-50-126.us-east-2.compute.internal
Pending
csr-c57lv 5m26s system:node:ip-10-0-95-157.us-east-2.compute.internal
Pending
...
```

- If the remaining CSRs are not approved, and are in the **Pending** status, approve the CSRs for your cluster machines:

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{"\n"}\n{{end}}{{end}}' | xargs oc adm certificate approve
```

- After all client and server CSRs have been approved, the machines have the **Ready** status. Verify this by running the following command:

```
$ oc get nodes
```

Example output

```
NAME      STATUS   ROLES    AGE   VERSION
master-0  Ready   master   73m   v1.29.4
master-1  Ready   master   73m   v1.29.4
master-2  Ready   master   74m   v1.29.4
worker-0  Ready   worker   11m   v1.29.4
worker-1  Ready   worker   11m   v1.29.4
```



NOTE

It can take a few minutes after approval of the server CSRs for the machines to transition to the **Ready** status.

Additional information

- For more information on CSRs, see [Certificate Signing Requests](#).

15.4.15. Initial Operator configuration

After the control plane initializes, you must immediately configure some Operators so that they all become available.

Prerequisites

- Your control plane has initialized.

Procedure

1. Watch the cluster components come online:

```
$ watch -n5 oc get clusteroperators
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.16.0	True	False	False	19m
baremetal	4.16.0	True	False	False	37m
cloud-credential	4.16.0	True	False	False	40m
cluster-autoscaler	4.16.0	True	False	False	37m
config-operator	4.16.0	True	False	False	38m
console	4.16.0	True	False	False	26m
csi-snapshot-controller	4.16.0	True	False	False	37m
dns	4.16.0	True	False	False	37m
etcd	4.16.0	True	False	False	36m
image-registry	4.16.0	True	False	False	31m
ingress	4.16.0	True	False	False	30m
insights	4.16.0	True	False	False	31m
kube-apiserver	4.16.0	True	False	False	26m
kube-controller-manager	4.16.0	True	False	False	36m
kube-scheduler	4.16.0	True	False	False	36m
kube-storage-version-migrator	4.16.0	True	False	False	37m
machine-api	4.16.0	True	False	False	29m
machine-approver	4.16.0	True	False	False	37m
machine-config	4.16.0	True	False	False	36m
marketplace	4.16.0	True	False	False	37m
monitoring	4.16.0	True	False	False	29m
network	4.16.0	True	False	False	38m
node-tuning	4.16.0	True	False	False	37m
openshift-apiserver	4.16.0	True	False	False	32m
openshift-controller-manager	4.16.0	True	False	False	30m
openshift-samples	4.16.0	True	False	False	32m
operator-lifecycle-manager	4.16.0	True	False	False	37m
operator-lifecycle-manager-catalog	4.16.0	True	False	False	37m
operator-lifecycle-manager-packageserver	4.16.0	True	False	False	32m
service-ca	4.16.0	True	False	False	38m
storage	4.16.0	True	False	False	37m

2. Configure the Operators that are not available.

Additional resources

- See [Gathering logs from a failed installation](#) for details about gathering data in the event of a failed OpenShift Container Platform installation.
- See [Troubleshooting Operator issues](#) for steps to check Operator pod health across the cluster and gather Operator logs for diagnosis.

15.4.15.1. Disabling the default OperatorHub catalog sources

Operator catalogs that source content provided by Red Hat and community projects are configured for OperatorHub by default during an OpenShift Container Platform installation. In a restricted network environment, you must disable the default catalogs as a cluster administrator.

Procedure

- Disable the sources for the default catalogs by adding **disableAllDefaultSources: true** to the **OperatorHub** object:

```
$ oc patch OperatorHub cluster --type json \
  -p '[{"op": "add", "path": "/spec/disableAllDefaultSources", "value": true}]'
```

TIP

Alternatively, you can use the web console to manage catalog sources. From the **Administration** → **Cluster Settings** → **Configuration** → **OperatorHub** page, click the **Sources** tab, where you can create, update, delete, disable, and enable individual sources.

15.4.15.2. Image registry storage configuration

The Image Registry Operator is not initially available for platforms that do not provide default storage. After installation, you must configure your registry to use storage so that the Registry Operator is made available.

Instructions are shown for configuring a persistent volume, which is required for production clusters. Where applicable, instructions are shown for configuring an empty directory as the storage location, which is available for only non-production clusters.

Additional instructions are provided for allowing the image registry to use block storage types by using the **Recreate** rollout strategy during upgrades.

15.4.15.2.1. Changing the image registry's management state

To start the image registry, you must change the Image Registry Operator configuration's **managementState** from **Removed** to **Managed**.

Procedure

- Change **managementState** Image Registry Operator configuration from **Removed** to **Managed**. For example:

```
$ oc patch configs.imageregistry.operator.openshift.io cluster --type merge --patch '{"spec": {"managementState": "Managed"}}'
```

15.4.15.2.2. Configuring registry storage for bare metal and other manual installations

As a cluster administrator, following installation you must configure your registry to use storage.

Prerequisites

- You have access to the cluster as a user with the **cluster-admin** role.
- You have a cluster that uses manually-provisioned Red Hat Enterprise Linux CoreOS (RHCOS) nodes, such as bare metal.
- You have provisioned persistent storage for your cluster, such as Red Hat OpenShift Data Foundation.



IMPORTANT

OpenShift Container Platform supports **ReadWriteOnce** access for image registry storage when you have only one replica. **ReadWriteOnce** access also requires that the registry uses the **Recreate** rollout strategy. To deploy an image registry that supports high availability with two or more replicas, **ReadWriteMany** access is required.

- Must have 100Gi capacity.

Procedure

1. To configure your registry to use storage, change the **spec.storage.pvc** in the **configs.imageregistry/cluster** resource.



NOTE

When you use shared storage, review your security settings to prevent outside access.

2. Verify that you do not have a registry pod:

```
$ oc get pod -n openshift-image-registry -l docker-registry=default
```

Example output

```
No resources found in openshift-image-registry namespace
```



NOTE

If you do have a registry pod in your output, you do not need to continue with this procedure.

3. Check the registry configuration:

```
$ oc edit configs.imageregistry.operator.openshift.io
```

Example output

```
storage:
  pvc:
    claim:
```

Leave the **claim** field blank to allow the automatic creation of an **image-registry-storage** PVC.

4. Check the **clusteroperator** status:

```
$ oc get clusteroperator image-registry
```

Example output

```
NAME          VERSION          AVAILABLE PROGRESSING DEGRADED SINCE
MESSAGE
image-registry 4.16             True      False      False      6h50m
```

5. Ensure that your registry is set to managed to enable building and pushing of images.

- Run:

```
$ oc edit configs.imageregistry/cluster
```

Then, change the line

```
managementState: Removed
```

to

```
managementState: Managed
```

15.4.15.2.3. Configuring storage for the image registry in non-production clusters

You must configure storage for the Image Registry Operator. For non-production clusters, you can set the image registry to an empty directory. If you do so, all images are lost if you restart the registry.

Procedure

- To set the image registry storage to an empty directory:

```
$ oc patch configs.imageregistry.operator.openshift.io cluster --type merge --patch '{"spec": {"storage":{"emptyDir":{}}}'
```



WARNING

Configure this option for only non-production clusters.


If you run this command before the Image Registry Operator initializes its components, the **oc patch** command fails with the following error:

```
Error from server (NotFound): configs.imageregistry.operator.openshift.io "cluster" not found
```

Wait a few minutes and run the command again.

15.4.15.2.4. Configuring block registry storage for bare metal

To allow the image registry to use block storage types during upgrades as a cluster administrator, you can use the **Recreate** rollout strategy.



IMPORTANT

Block storage volumes, or block persistent volumes, are supported but not recommended for use with the image registry on production clusters. An installation where the registry is configured on block storage is not highly available because the registry cannot have more than one replica.

If you choose to use a block storage volume with the image registry, you must use a filesystem persistent volume claim (PVC).

Procedure

1. Enter the following command to set the image registry storage as a block storage type, patch the registry so that it uses the **Recreate** rollout strategy, and runs with only one (**1**) replica:

```
$ oc patch config.imageregistry.operator.openshift.io/cluster --type=merge -p '{"spec": {"rolloutStrategy": "Recreate", "replicas": 1}}'
```

2. Provision the PV for the block storage device, and create a PVC for that volume. The requested block volume uses the ReadWriteOnce (RWO) access mode.
 - a. Create a **pvc.yaml** file with the following contents to define a VMware vSphere **PersistentVolumeClaim** object:

```
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: image-registry-storage 1
  namespace: openshift-image-registry 2
spec:
  accessModes:
  - ReadWriteOnce 3
  resources:
    requests:
      storage: 100Gi 4
```

- 1** A unique name that represents the **PersistentVolumeClaim** object.
- 2** The namespace for the **PersistentVolumeClaim** object, which is **openshift-image-registry**.
- 3** The access mode of the persistent volume claim. With **ReadWriteOnce**, the volume can be mounted with read and write permissions by a single node.

- 4 The size of the persistent volume claim.

- b. Enter the following command to create the **PersistentVolumeClaim** object from the file:

```
$ oc create -f pvc.yaml -n openshift-image-registry
```

3. Enter the following command to edit the registry configuration so that it references the correct PVC:

```
$ oc edit config.imageregistry.operator.openshift.io -o yaml
```

Example output

```
storage:
  pvc:
    claim: 1
```

- 1 By creating a custom PVC, you can leave the **claim** field blank for the default automatic creation of an **image-registry-storage** PVC.

15.4.16. Completing installation on user-provisioned infrastructure

After you complete the Operator configuration, you can finish installing the cluster on infrastructure that you provide.

Prerequisites

- Your control plane has initialized.
- You have completed the initial Operator configuration.

Procedure

1. Confirm that all the cluster components are online with the following command:

```
$ watch -n5 oc get clusteroperators
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.16.0	True	False	False	19m
baremetal	4.16.0	True	False	False	37m
cloud-credential	4.16.0	True	False	False	40m
cluster-autoscaler	4.16.0	True	False	False	37m
config-operator	4.16.0	True	False	False	38m
console	4.16.0	True	False	False	26m
csi-snapshot-controller	4.16.0	True	False	False	37m
dns	4.16.0	True	False	False	37m
etcd	4.16.0	True	False	False	36m

image-registry	4.16.0	True	False	False	31m
ingress	4.16.0	True	False	False	30m
insights	4.16.0	True	False	False	31m
kube-apiserver	4.16.0	True	False	False	26m
kube-controller-manager	4.16.0	True	False	False	36m
kube-scheduler	4.16.0	True	False	False	36m
kube-storage-version-migrator	4.16.0	True	False	False	37m
machine-api	4.16.0	True	False	False	29m
machine-approver	4.16.0	True	False	False	37m
machine-config	4.16.0	True	False	False	36m
marketplace	4.16.0	True	False	False	37m
monitoring	4.16.0	True	False	False	29m
network	4.16.0	True	False	False	38m
node-tuning	4.16.0	True	False	False	37m
openshift-apiserver	4.16.0	True	False	False	32m
openshift-controller-manager	4.16.0	True	False	False	30m
openshift-samples	4.16.0	True	False	False	32m
operator-lifecycle-manager	4.16.0	True	False	False	37m
operator-lifecycle-manager-catalog	4.16.0	True	False	False	37m
operator-lifecycle-manager-packageserver	4.16.0	True	False	False	32m
service-ca	4.16.0	True	False	False	38m
storage	4.16.0	True	False	False	37m

Alternatively, the following command notifies you when all of the clusters are available. It also retrieves and displays credentials:


```
$ ./openshift-install --dir <installation_directory> wait-for install-complete 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

Example output

```
INFO Waiting up to 30m0s for the cluster to initialize...
```

The command succeeds when the Cluster Version Operator finishes deploying the OpenShift Container Platform cluster from Kubernetes API server.



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

2. Confirm that the Kubernetes API server is communicating with the pods.

a. To view a list of all pods, use the following command:

```
$ oc get pods --all-namespaces
```

Example output

```

NAMESPACE           NAME                                     READY  STATUS
RESTARTS  AGE
openshift-apiserver-operator  openshift-apiserver-operator-85cb746d55-zqhs8  1/1
Running  1      9m
openshift-apiserver          apiserver-67b9g                                1/1  Running  0
3m
openshift-apiserver          apiserver-ljcmx                                1/1  Running  0
1m
openshift-apiserver          apiserver-z25h4                                1/1  Running  0
2m
openshift-authentication-operator  authentication-operator-69d5d8bf84-vh2n8      1/1
Running  0      5m
...

```

b. View the logs for a pod that is listed in the output of the previous command by using the following command:

```
$ oc logs <pod_name> -n <namespace> 1
```

- 1** Specify the pod name and namespace, as shown in the output of the previous command.

If the pod logs display, the Kubernetes API server can communicate with the cluster machines.

3. For an installation with Fibre Channel Protocol (FCP), additional steps are required to enable multipathing. Do not enable multipathing during installation. See "Enabling multipathing with kernel arguments on RHCOS" in the *Postinstallation machine configuration tasks* documentation for more information.

4. Register your cluster on the [Cluster registration](#) page.

15.4.17. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service

15.4.18. Next steps

- [Validating an installation](#).
- [Customize your cluster](#).
- [Configure image streams](#) for the Cluster Samples Operator and the **must-gather** tool.
- Learn how to [use Operator Lifecycle Manager \(OLM\) on restricted networks](#).
- If the mirror registry that you used to install your cluster has a trusted CA, add it to the cluster by [configuring additional trust stores](#).
- If necessary, you can [opt out of remote health reporting](#).
- If necessary, see [Registering your disconnected cluster](#)

15.5. SCALING A USER-PROVISIONED CLUSTER WITH THE BARE METAL OPERATOR

After deploying a user-provisioned infrastructure cluster, you can use the Bare Metal Operator (BMO) and other metal3 components to scale bare-metal hosts in the cluster. This approach helps you to scale a user-provisioned cluster in a more automated way.

15.5.1. About scaling a user-provisioned cluster with the Bare Metal Operator

You can scale user-provisioned infrastructure clusters by using the Bare Metal Operator (BMO) and other metal3 components. User-provisioned infrastructure installations do not feature the Machine API Operator. The Machine API Operator typically manages the lifecycle of bare-metal hosts in a cluster. However, it is possible to use the BMO and other metal3 components to scale nodes in user-provisioned clusters without requiring the Machine API Operator.

15.5.1.1. Prerequisites for scaling a user-provisioned cluster

- You installed a user-provisioned infrastructure cluster on bare metal.
- You have baseboard management controller (BMC) access to the hosts.

15.5.1.2. Limitations for scaling a user-provisioned cluster

- You cannot use a provisioning network to scale user-provisioned infrastructure clusters by using the Bare Metal Operator (BMO).
 - Consequentially, you can only use bare-metal host drivers that support virtual media networking booting, for example **redfish-virtualmedia** and **idrac-virtualmedia**.
- You cannot scale **MachineSet** objects in user-provisioned infrastructure clusters by using the BMO.

15.5.2. Configuring a provisioning resource to scale user-provisioned clusters

Create a **Provisioning** custom resource (CR) to enable Metal platform components on a user-provisioned infrastructure cluster.

Prerequisites

- You installed a user-provisioned infrastructure cluster on bare metal.

Procedure

1. Create a **Provisioning** CR.
 - a. Save the following YAML in the **provisioning.yaml** file:

```
apiVersion: metal3.io/v1alpha1
kind: Provisioning
metadata:
  name: provisioning-configuration
spec:
  provisioningNetwork: "Disabled"
  watchAllNamespaces: false
```



NOTE

OpenShift Container Platform 4.16 does not support enabling a provisioning network when you scale a user-provisioned cluster by using the Bare Metal Operator.

2. Create the **Provisioning** CR by running the following command:

```
$ oc create -f provisioning.yaml
```

Example output

```
provisioning.metal3.io/provisioning-configuration created
```

Verification

- Verify that the provisioning service is running by running the following command:

```
$ oc get pods -n openshift-machine-api
```

Example output

NAME	READY	STATUS	RESTARTS	AGE
cluster-autoscaler-operator-678c476f4c-jjdn5	2/2	Running	0	5d21h
cluster-baremetal-operator-6866f7b976-gmvgh	2/2	Running	0	5d21h
control-plane-machine-set-operator-7d8566696c-bh4jz	1/1	Running	0	5d21h
ironic-proxy-64bdw	1/1	Running	0	5d21h
ironic-proxy-rbggf	1/1	Running	0	5d21h
ironic-proxy-vj54c	1/1	Running	0	5d21h
machine-api-controllers-544d6849d5-tgj9l	7/7	Running	1 (5d21h ago)	5d21h
machine-api-operator-5c4ff4b86d-6fjmq	2/2	Running	0	5d21h
metal3-6d98f84cc8-zn2mx	5/5	Running	0	5d21h
metal3-image-customization-59d745768d-bhrp7	1/1	Running	0	5d21h

15.5.3. Provisioning new hosts in a user-provisioned cluster by using the BMO

You can use the Bare Metal Operator (BMO) to provision bare-metal hosts in a user-provisioned cluster by creating a **BareMetalHost** custom resource (CR).



NOTE

To provision bare-metal hosts to the cluster by using the BMO, you must set the **spec.externallyProvisioned** specification in the **BareMetalHost** custom resource to **false**.

Prerequisites

- You created a user-provisioned bare-metal cluster.
- You have baseboard management controller (BMC) access to the hosts.
- You deployed a provisioning service in the cluster by creating a **Provisioning** CR.

Procedure

1. Create the **Secret** CR and the **BareMetalHost** CR.
 - a. Save the following YAML in the **bmh.yaml** file:

```
---
apiVersion: v1
kind: Secret
metadata:
  name: worker1-bmc
  namespace: openshift-machine-api
type: Opaque
data:
  username: <base64_of_uid>
  password: <base64_of_pwd>
---
apiVersion: metal3.io/v1alpha1
kind: BareMetalHost
metadata:
  name: worker1
```

```

namespace: openshift-machine-api
spec:
  bmc:
    address: <protocol>://<bmc_url> 1
    credentialsName: "worker1-bmc"
  bootMACAddress: <nic1_mac_address>
  externallyProvisioned: false 2
  customDeploy:
    method: install_coreos
  online: true
  userData:
    name: worker-user-data-managed
    namespace: openshift-machine-api

```

- 1** You can only use bare-metal host drivers that support virtual media networking booting, for example **redfish-virtualmedia** and **idrac-virtualmedia**.
- 2** You must set the **spec.externallyProvisioned** specification in the **BareMetalHost** custom resource to **false**. The default value is **false**.

2. Create the bare-metal host object by running the following command:

```
$ oc create -f bmh.yaml
```

Example output

```
secret/worker1-bmc created
baremetalhost.metal3.io/worker1 created
```

3. Approve all certificate signing requests (CSRs).
 - a. Verify that the provisioning state of the host is **provisioned** by running the following command:

```
$ oc get bmh -A
```

Example output

NAMESPACE	NAME	STATE	CONSUMER	ONLINE
ERROR	AGE			
openshift-machine-api	controller1	externally provisioned	true	5m25s
openshift-machine-api	worker1	provisioned	true	4m45s

- b. Get the list of pending CSRs by running the following command:

```
$ oc get csr
```

Example output

NAME	AGE	SIGNERNAME	REQUESTOR
REQUESTEDDURATION	CONDITION		
csr-gfm9f	33s	kubernetes.io/kube-apiserver-client-kubelet	

```
system:serviceaccount:openshift-machine-config-operator:node-bootstrapper <none> Pending
```

- c. Approve the CSR by running the following command:

```
$ oc adm certificate approve <csr_name>
```

Example output

```
certificatesigningrequest.certificates.k8s.io/<csr_name> approved
```

Verification

- Verify that the node is ready by running the following command:

```
$ oc get nodes
```

Example output

```
NAME          STATUS  ROLES    AGE   VERSION
app1          Ready   worker   47s   v1.24.0+dc5a2fd
controller1   Ready   master,worker 2d22h v1.24.0+dc5a2fd
```

15.5.4. Optional: Managing existing hosts in a user-provisioned cluster by using the BMO

Optionally, you can use the Bare Metal Operator (BMO) to manage existing bare-metal controller hosts in a user-provisioned cluster by creating a **BareMetalHost** object for the existing host. It is not a requirement to manage existing user-provisioned hosts; however, you can enroll them as externally-provisioned hosts for inventory purposes.



IMPORTANT

To manage existing hosts by using the BMO, you must set the **spec.externallyProvisioned** specification in the **BareMetalHost** custom resource to **true** to prevent the BMO from re-provisioning the host.

Prerequisites

- You created a user-provisioned bare-metal cluster.
- You have baseboard management controller (BMC) access to the hosts.
- You deployed a provisioning service in the cluster by creating a **Provisioning** CR.

Procedure

1. Create the **Secret** CR and the **BareMetalHost** CR.

- a. Save the following YAML in the **controller.yaml** file:

```
---
```

```

apiVersion: v1
kind: Secret
metadata:
  name: controller1-bmc
  namespace: openshift-machine-api
type: Opaque
data:
  username: <base64_of_uid>
  password: <base64_of_pwd>
---
apiVersion: metal3.io/v1alpha1
kind: BareMetalHost
metadata:
  name: controller1
  namespace: openshift-machine-api
spec:
  bmc:
    address: <protocol>://<bmc_url> 1
    credentialsName: "controller1-bmc"
    bootMACAddress: <nic1_mac_address>
  customDeploy:
    method: install_coreos
  externallyProvisioned: true 2
  online: true
  userData:
    name: controller-user-data-managed
    namespace: openshift-machine-api

```

- 1** You can only use bare-metal host drivers that support virtual media networking booting, for example **redfish-virtualmedia** and **idrac-virtualmedia**.
- 2** You must set the value to true to prevent the BMO from re-provisioning the bare-metal controller host.

2. Create the bare-metal host object by running the following command:

```
$ oc create -f controller.yaml
```

Example output

```
secret/controller1-bmc created
baremetalhost.metal3.io/controller1 created
```

Verification

- Verify that the BMO created the bare-metal host object by running the following command:

```
$ oc get bmh -A
```

Example output

NAMESPACE	NAME	STATE	CONSUMER	ONLINE	ERROR
openshift-machine-api	controller1	externally provisioned		true	13s

15.5.5. Removing hosts from a user-provisioned cluster by using the BMO

You can use the Bare Metal Operator (BMO) to remove bare-metal hosts from a user-provisioned cluster.

Prerequisites

- You created a user-provisioned bare-metal cluster.
- You have baseboard management controller (BMC) access to the hosts.
- You deployed a provisioning service in the cluster by creating a **Provisioning** CR.

Procedure

1. Cordon and drain the host by running the following command:

```
$ oc adm drain app1 --force --ignore-daemonsets=true
```

Example output

```
node/app1 cordoned
WARNING: ignoring DaemonSet-managed Pods: openshift-cluster-node-tuning-
operator/tuned-tvthg, openshift-dns/dns-
default-9q6rz, openshift-dns/node-resolver-zvt42, openshift-image-registry/node-ca-mzxt,
openshift-ingress-cana
ry/ingress-canary-qq5lf, openshift-machine-config-operator/machine-config-daemon-v79dm,
openshift-monitoring/nod
e-exporter-2vn59, openshift-multus/multus-additional-cni-plugins-wssvj, openshift-
multus/multus-fn8tg, openshift
-multus/network-metrics-daemon-5qv55, openshift-network-diagnostics/network-check-
target-jqxn2, openshift-ovn-ku
bernetes/ovnkube-node-rsvqg
evicting pod openshift-operator-lifecycle-manager/collect-profiles-27766965-258vp
evicting pod openshift-operator-lifecycle-manager/collect-profiles-27766950-kg5mk
evicting pod openshift-operator-lifecycle-manager/collect-profiles-27766935-stf4s
pod/collect-profiles-27766965-258vp evicted
pod/collect-profiles-27766950-kg5mk evicted
pod/collect-profiles-27766935-stf4s evicted
node/app1 drained
```

2. Delete the **customDeploy** specification from the **BareMetalHost** CR.
 - a. Edit the **BareMetalHost** CR for the host by running the following command:

```
$ oc edit bmh -n openshift-machine-api <host_name>
```

- b. Delete the lines **spec.customDeploy** and **spec.customDeploy.method**:

```
...
  customDeploy:
    method: install_coreos
```

- c. Verify that the provisioning state of the host changes to **deprovisioning** by running the following command:

```
$ oc get bmh -A
```

Example output

NAMESPACE	NAME	STATE	CONSUMER	ONLINE
ERROR	AGE			
openshift-machine-api	controller1	externally provisioned	true	58m
openshift-machine-api	worker1	deprovisioning	true	57m

3. Delete the node by running the following command:

```
$ oc delete node <node_name>
```

Verification

- Verify the node is deleted by running the following command:

```
$ oc get nodes
```

Example output


NAME	STATUS	ROLES	AGE	VERSION
controller1	Ready	master,worker	2d23h	v1.24.0+dc5a2fd

15.6. INSTALLATION CONFIGURATION PARAMETERS FOR BARE METAL

Before you deploy an OpenShift Container Platform cluster, you provide a customized **install-config.yaml** installation configuration file that describes the details for your environment.

15.6.1. Available installation configuration parameters for bare metal

The following tables specify the required, optional, and bare metal-specific installation configuration parameters that you can set as part of the installation process.



NOTE

After installation, you cannot modify these parameters in the **install-config.yaml** file.

15.6.1.1. Required configuration parameters

Required installation configuration parameters are described in the following table:

Table 15.42. Required parameters

Parameter	Description	Values
apiVersion:	The API version for the install-config.yaml content. The current version is v1 . The installation program may also support older API versions.	String
baseDomain:	The base domain of your cloud provider. The base domain is used to create routes to your OpenShift Container Platform cluster components. The full DNS name for your cluster is a combination of the baseDomain and metadata.name parameter values that uses the <metadata.name> . <baseDomain> format.	A fully-qualified domain or subdomain name, such as example.com .
metadata:	Kubernetes resource ObjectMeta , from which only the name parameter is consumed.	Object
metadata: name:	The name of the cluster. DNS records for the cluster are all subdomains of {{.metadata.name}} . {{.baseDomain}} .	String of lowercase letters and hyphens (-), such as dev .
platform:	The configuration for the specific platform upon which to perform the installation: aws , baremetal , azure , gcp , ibmcloud , nutanix , openstack , powervs , vsphere , or {} . For additional information about platform . <platform> parameters, consult the table for your specific platform that follows.	Object

Parameter	Description	Values
<code>pullSecret:</code>	Get a pull secret from Red Hat OpenShift Cluster Manager to authenticate downloading container images for OpenShift Container Platform components from services such as Quay.io.	<pre>{ "auths":{ "cloud.openshift.com":{ "auth":"b3Blb=", "email":"you@example.com" }, "quay.io":{ "auth":"b3Blb=", "email":"you@example.com" } } }</pre>

15.6.1.2. Network configuration parameters

You can customize your installation configuration based on the requirements of your existing network infrastructure. For example, you can expand the IP address block for the cluster network or provide different IP address blocks than the defaults.

- If you use the Red Hat OpenShift Networking OVN-Kubernetes network plugin, both IPv4 and IPv6 address families are supported.

If you configure your cluster to use both IP address families, review the following requirements:

- Both IP families must use the same network interface for the default gateway.
- Both IP families must have the default gateway.
- You must specify IPv4 and IPv6 addresses in the same order for all network configuration parameters. For example, in the following configuration IPv4 addresses are listed before IPv6 addresses.


```
networking:
  clusterNetwork:
    - cidr: 10.128.0.0/14
      hostPrefix: 23
    - cidr: fd00:10:128::/56
      hostPrefix: 64
  serviceNetwork:
    - 172.30.0.0/16
    - fd00:172:16::/112
```




NOTE

Globalnet is not supported with Red Hat OpenShift Data Foundation disaster recovery solutions. For regional disaster recovery scenarios, ensure that you use a nonoverlapping range of private IP addresses for the cluster and service networks in each cluster.

Table 15.43. Network parameters

Parameter	Description	Values
<code>networking:</code>	The configuration for the cluster network.	Object  NOTE You cannot modify parameters specified by the networking object after installation.
<code>networking:</code> <code>networkType:</code>	The Red Hat OpenShift Networking network plugin to install.	OVNKubernetes. OVNKubernetes is a CNI plugin for Linux networks and hybrid networks that contain both Linux and Windows servers. The default value is OVNKubernetes .
<code>networking:</code> <code>clusterNetwork:</code>	The IP address blocks for pods. The default value is 10.128.0.0/14 with a host prefix of /23 . If you specify multiple IP address blocks, the blocks must not overlap.	An array of objects. For example: <code>networking:</code> <code>clusterNetwork:</code> - cidr: 10.128.0.0/14 <code>hostPrefix:</code> 23 - cidr: fd01::/48 <code>hostPrefix:</code> 64
<code>networking:</code> <code>clusterNetwork:</code> <code>cidr:</code>	Required if you use networking.clusterNetwork . An IP address block. If you use the OVN-Kubernetes network plugin, you can specify IPv4 and IPv6 networks.	An IP address block in Classless Inter-Domain Routing (CIDR) notation. The prefix length for an IPv4 block is between 0 and 32 . The prefix length for an IPv6 block is between 0 and 128 . For example, 10.128.0.0/14 or fd01::/48 .
<code>networking:</code> <code>clusterNetwork:</code> <code>hostPrefix:</code>	The subnet prefix length to assign to each individual node. For example, if hostPrefix is set to 23 then each node is assigned a /23 subnet out of the given cidr . A hostPrefix value of 23 provides 510 ($2^{(32 - 23)} - 2$) pod IP addresses.	A subnet prefix. For an IPv4 network the default value is 23 . For an IPv6 network the default value is 64 . The default value is also the minimum value for IPv6.

Parameter	Description	Values
<code>networking: serviceNetwork:</code>	<p>The IP address block for services. The default value is 172.30.0.0/16.</p> <p>The OVN-Kubernetes network plugins supports only a single IP address block for the service network.</p> <p>If you use the OVN-Kubernetes network plugin, you can specify an IP address block for both of the IPv4 and IPv6 address families.</p>	<p>An array with an IP address block in CIDR format. For example:</p> <pre>networking: serviceNetwork: - 172.30.0.0/16 - fd02::/112</pre>
<code>networking: machineNetwork:</code>	<p>The IP address blocks for machines.</p> <p>If you specify multiple IP address blocks, the blocks must not overlap.</p>	<p>An array of objects. For example:</p> <pre>networking: machineNetwork: - cidr: 10.0.0.0/16</pre>
<code>networking: machineNetwork: cidr:</code>	<p>Required if you use networking.machineNetwork. An IP address block. The default value is 10.0.0.0/16 for all platforms other than libvirt and IBM Power® Virtual Server. For libvirt, the default value is 192.168.126.0/24. For IBM Power® Virtual Server, the default value is 192.168.0.0/24.</p>	<p>An IP network block in CIDR notation.</p> <p>For example, 10.0.0.0/16 or fd00::/48.</p> <div style="display: flex; align-items: flex-start;">  <div> <p>NOTE</p> <p>Set the networking.machineNetwork to match the CIDR that the preferred NIC resides in.</p> </div> </div>


15.6.1.3. Optional configuration parameters


Optional installation configuration parameters are described in the following table:

Table 15.44. Optional parameters

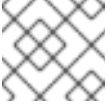

Parameter	Description	Values
<code>additionalTrustBundle:</code>	<p>A PEM-encoded X.509 certificate bundle that is added to the nodes' trusted certificate store. This trust bundle may also be used when a proxy has been configured.</p>	String


Parameter	Description	Values
capabilities:	Controls the installation of optional core cluster components. You can reduce the footprint of your OpenShift Container Platform cluster by disabling optional components. For more information, see the "Cluster capabilities" page in <i>Installing</i> .	String array
capabilities: baselineCapabilitySet:	Selects an initial set of optional capabilities to enable. Valid values are None , v4.11 , v4.12 and vCurrent . The default value is vCurrent .	String
capabilities: additionalEnabledCapabilities:	Extends the set of optional capabilities beyond what you specify in baselineCapabilitySet . You may specify multiple capabilities in this parameter.	String array
cpuPartitioningMode:	Enables workload partitioning, which isolates OpenShift Container Platform services, cluster management workloads, and infrastructure pods to run on a reserved set of CPUs. Workload partitioning can only be enabled during installation and cannot be disabled after installation. While this field enables workload partitioning, it does not configure workloads to use specific CPUs. For more information, see the <i>Workload partitioning</i> page in the <i>Scalability and Performance</i> section.	None or AllNodes . None is the default value.
compute:	The configuration for the machines that comprise the compute nodes.	Array of MachinePool objects.
compute: architecture:	Determines the instruction set architecture of the machines in the pool. Currently, clusters with varied architectures are not supported. All pools must specify the same architecture. Valid values are amd64 and arm64 .	String

Parameter	Description	Values
<code>compute: hyperthreading:</code>	<p>Whether to enable or disable simultaneous multithreading, or hyperthreading, on compute machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.</p>  <p>IMPORTANT</p> <p>If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.</p>	Enabled or Disabled
<code>compute: name:</code>	Required if you use compute . The name of the machine pool.	worker
<code>compute: platform:</code>	Required if you use compute . Use this parameter to specify the cloud provider to host the worker machines. This parameter value must match the controlPlane.platform parameter value.	aws, azure, gcp, ibmcloud, nutanix, openstack, powervs, vsphere , or <code>{}</code>
<code>compute: replicas:</code>	The number of compute machines, which are also known as worker machines, to provision.	A positive integer greater than or equal to 2 . The default value is 3 .
<code>featureSet:</code>	Enables the cluster for a feature set. A feature set is a collection of OpenShift Container Platform features that are not enabled by default. For more information about enabling a feature set during installation, see "Enabling features using feature gates".	String. The name of the feature set to enable, such as TechPreviewNoUpgrade .
<code>controlPlane:</code>	The configuration for the machines that comprise the control plane.	Array of MachinePool objects.

Parameter	Description	Values
<code>controlPlane: architecture:</code>	Determines the instruction set architecture of the machines in the pool. Currently, clusters with varied architectures are not supported. All pools must specify the same architecture. Valid values are amd64 and arm64 .	String
<code>controlPlane: hyperthreading:</code>	Whether to enable or disable simultaneous multithreading, or hyperthreading , on control plane machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.  IMPORTANT If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.	Enabled or Disabled
<code>controlPlane: name:</code>	Required if you use controlPlane . The name of the machine pool.	master
<code>controlPlane: platform:</code>	Required if you use controlPlane . Use this parameter to specify the cloud provider that hosts the control plane machines. This parameter value must match the compute.platform parameter value.	aws, azure, gcp, ibmcloud, nutanix, openstack, powervs, vsphere, or {}
<code>controlPlane: replicas:</code>	The number of control plane machines to provision.	Supported values are 3 , or 1 when deploying single-node OpenShift.

Parameter	Description	Values
credentialsMode:	<p>The Cloud Credential Operator (CCO) mode. If no mode is specified, the CCO dynamically tries to determine the capabilities of the provided credentials, with a preference for mint mode on the platforms where multiple modes are supported.</p>	Mint, Passthrough, Manual or an empty string (""). ^[1]
fips:	<p>Enable or disable FIPS mode. The default is false (disabled). If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.</p> <div data-bbox="486 884 593 1937" style="background-color: black; color: white; padding: 10px; margin: 10px 0;"> <p>IMPORTANT</p> <p>To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see Installing the system in FIPS mode.</p> <p>When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.</p> </div>	false or true

Parameter	Description	NOTE	Values
		<p>If you are using Azure File storage, you cannot enable FIPS mode.</p>	
imageContentSources:	<p>Sources and repositories for the release-image content.</p>		<p>Array of objects. Includes a source and, optionally, mirrors, as described in the following rows of this table.</p>
imageContentSources: source:	<p>Required if you use imageContentSources. Specify the repository that users refer to, for example, in image pull specifications.</p>		<p>String</p>
imageContentSources: mirrors:	<p>Specify one or more repositories that may also contain the same images.</p>		<p>Array of strings</p>
publish:	<p>How to publish or expose the user-facing endpoints of your cluster, such as the Kubernetes API, OpenShift routes.</p>	<p>Internal or External. The default value is External.</p> <p>Setting this field to Internal is not supported on non-cloud platforms.</p> <div data-bbox="979 1496 1083 1749" style="display: inline-block; vertical-align: top;">  </div> <p>IMPORTANT</p> <p>If the value of the field is set to Internal, the cluster will become non-functional. For more information, refer to BZ#1953035.</p>	

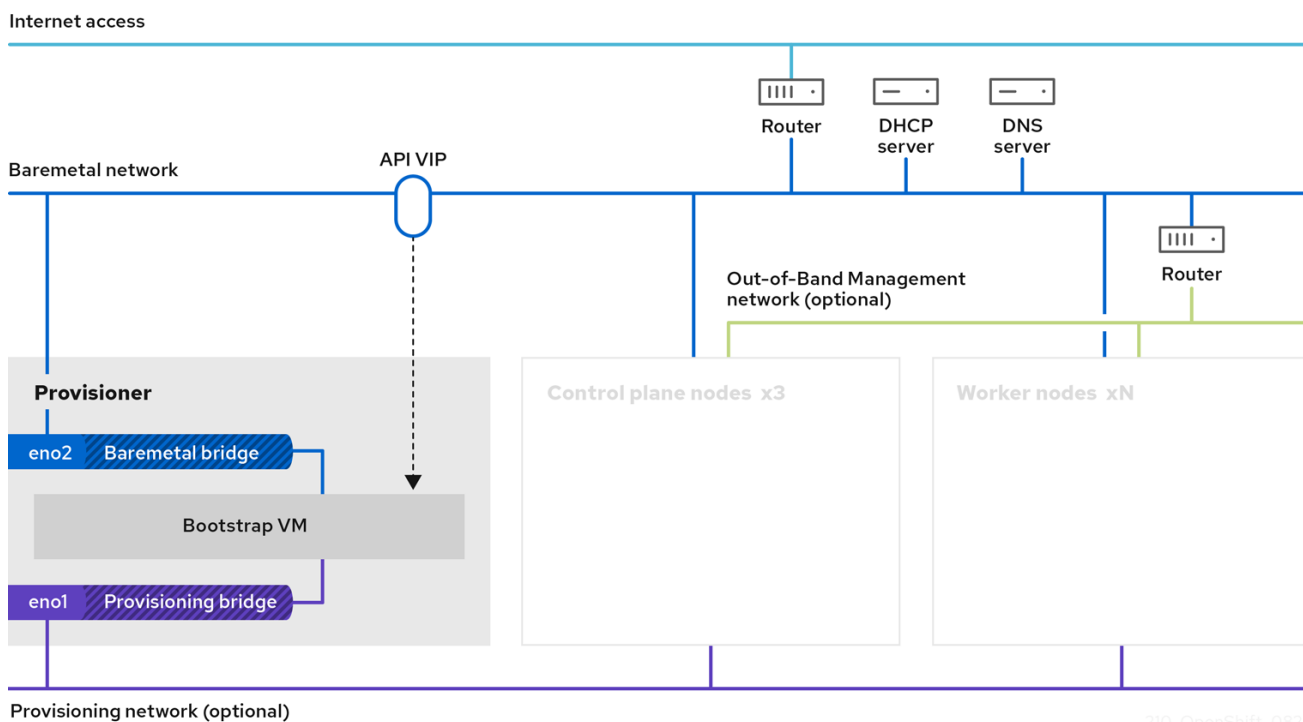
Parameter	Description	Values
sshKey:	<p>The SSH key to authenticate access to your cluster machines.</p>  <p>NOTE</p> <p>For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your ssh-agent process uses.</p>	<p>For example, sshKey: ssh-ed25519 AAAA...</p>

1. Not all CCO modes are supported for all cloud providers. For more information about CCO modes, see the "Managing cloud provider credentials" entry in the *Authentication and authorization* content.

CHAPTER 16. DEPLOYING INSTALLER-PROVISIONED CLUSTERS ON BARE METAL

16.1. OVERVIEW

Installer-provisioned installation on bare metal nodes deploys and configures the infrastructure that an OpenShift Container Platform cluster runs on. This guide provides a methodology to achieving a successful installer-provisioned bare-metal installation. The following diagram illustrates the installation environment in phase 1 of deployment:



For the installation, the key elements in the previous diagram are:

- **Provisioner:** A physical machine that runs the installation program and hosts the bootstrap VM that deploys the control plane of a new OpenShift Container Platform cluster.
- **Bootstrap VM:** A virtual machine used in the process of deploying an OpenShift Container Platform cluster.
- **Network bridges:** The bootstrap VM connects to the bare metal network and to the provisioning network, if present, via network bridges, **eno1** and **eno2**.
- **API VIP:** An API virtual IP address (VIP) is used to provide failover of the API server across the control plane nodes. The API VIP first resides on the bootstrap VM. A script generates the **keepalived.conf** configuration file before launching the service. The VIP moves to one of the control plane nodes after the bootstrap process has completed and the bootstrap VM stops.

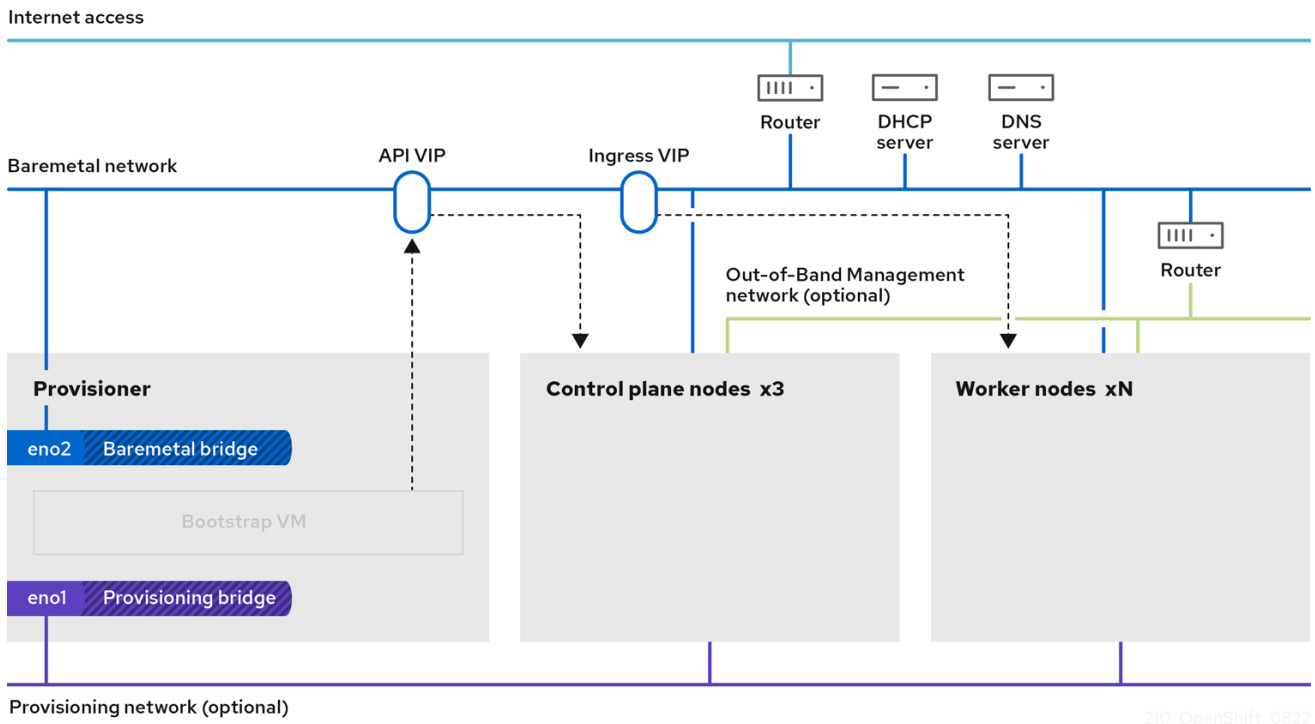
In phase 2 of the deployment, the provisioner destroys the bootstrap VM automatically and moves the virtual IP addresses (VIPs) to the appropriate nodes.

The **keepalived.conf** file sets the control plane machines with a lower Virtual Router Redundancy Protocol (VRRP) priority than the bootstrap VM, which ensures that the API on the control plane machines is fully functional before the API VIP moves from the bootstrap VM to the control plane. Once

the API VIP moves to one of the control plane nodes, traffic sent from external clients to the API VIP routes to an **haproxy** load balancer running on that control plane node. This instance of **haproxy** load balancer balances the API VIP traffic across the control plane nodes.

The Ingress VIP moves to the compute nodes. The **keepalived** instance also manages the Ingress VIP.

The following diagram illustrates phase 2 of deployment:



After this point, the node used by the provisioner can be removed or repurposed. From here, all additional provisioning tasks are carried out by the control plane.




NOTE

For installer-provisioned infrastructure installations, CoreDNS exposes port 53 at the node level, making it accessible from other routable networks.

Additional resources

- [Using DNS forwarding](#)



IMPORTANT

The provisioning network is optional, but it is required for PXE booting. If you deploy without a provisioning network, you must use a virtual media baseboard management controller (BMC) addressing option such as **redfish-virtualmedia** or **idrac-virtualmedia**.

16.2. PREREQUISITES

Installer-provisioned installation of OpenShift Container Platform requires:

1. One provisioner node with Red Hat Enterprise Linux (RHEL) 9.x installed. The provisioner can be removed after installation.

2. Three control plane nodes
3. Baseboard management controller (BMC) access to each node
4. At least one network:
 - a. One required routable network
 - b. One optional provisioning network
 - c. One optional management network

Before starting an installer-provisioned installation of OpenShift Container Platform, ensure the hardware environment meets the following requirements.

16.2.1. Node requirements

Installer-provisioned installation involves a number of hardware node requirements:

- **CPU architecture:** All nodes must use **x86_64** or **aarch64** CPU architecture.
- **Similar nodes:** Red Hat recommends nodes have an identical configuration per role. That is, Red Hat recommends nodes be the same brand and model with the same CPU, memory, and storage configuration.
- **Baseboard Management Controller:** The **provisioner** node must be able to access the baseboard management controller (BMC) of each OpenShift Container Platform cluster node. You may use IPMI, Redfish, or a proprietary protocol.
- **Latest generation:** Nodes must be of the most recent generation. Installer-provisioned installation relies on BMC protocols, which must be compatible across nodes. Additionally, RHEL 9.x ships with the most recent drivers for RAID controllers. Ensure that the nodes are recent enough to support RHEL 9.x for the **provisioner** node and RHCOS 9.x for the control plane and worker nodes.
- **Registry node:** (Optional) If setting up a disconnected mirrored registry, it is recommended the registry reside in its own node.
- **Provisioner node:** Installer-provisioned installation requires one **provisioner** node.
- **Control plane:** Installer-provisioned installation requires three control plane nodes for high availability. You can deploy an OpenShift Container Platform cluster with only three control plane nodes, making the control plane nodes schedulable as worker nodes. Smaller clusters are more resource efficient for administrators and developers during development, production, and testing.
- **Worker nodes:** While not required, a typical production cluster has two or more worker nodes.



IMPORTANT

Do not deploy a cluster with only one worker node, because the cluster will deploy with routers and ingress traffic in a degraded state.

- **Network interfaces:** Each node must have at least one network interface for the routable **baremetal** network. Each node must have one network interface for a **provisioning** network when using the **provisioning** network for deployment. Using the **provisioning** network is the


default configuration.



NOTE

Only one network card (NIC) on the same subnet can route traffic through the gateway. By default, Address Resolution Protocol (ARP) uses the lowest numbered NIC. Use a single NIC for each node in the same subnet to ensure that network load balancing works as expected. When using multiple NICs for a node in the same subnet, use a single bond or team interface. Then add the other IP addresses to that interface in the form of an alias IP address. If you require fault tolerance or load balancing at the network interface level, use an alias IP address on the bond or team interface. Alternatively, you can disable a secondary NIC on the same subnet or ensure that it has no IP address.

- **Unified Extensible Firmware Interface (UEFI):** Installer-provisioned installation requires UEFI boot on all OpenShift Container Platform nodes when using IPv6 addressing on the **provisioning** network. In addition, UEFI Device PXE Settings must be set to use the IPv6 protocol on the **provisioning** network NIC, but omitting the **provisioning** network removes this requirement.



IMPORTANT

When starting the installation from virtual media such as an ISO image, delete all old UEFI boot table entries. If the boot table includes entries that are not generic entries provided by the firmware, the installation might fail.

- **Secure Boot:** Many production scenarios require nodes with Secure Boot enabled to verify the node only boots with trusted software, such as UEFI firmware drivers, EFI applications, and the operating system. You may deploy with Secure Boot manually or managed.
 1. **Manually:** To deploy an OpenShift Container Platform cluster with Secure Boot manually, you must enable UEFI boot mode and Secure Boot on each control plane node and each worker node. Red Hat supports Secure Boot with manually enabled UEFI and Secure Boot only when installer-provisioned installations use Redfish virtual media. See "Configuring nodes for Secure Boot manually" in the "Configuring nodes" section for additional details.
 2. **Managed:** To deploy an OpenShift Container Platform cluster with managed Secure Boot, you must set the **bootMode** value to **UEFISecureBoot** in the **install-config.yaml** file. Red Hat only supports installer-provisioned installation with managed Secure Boot on 10th generation HPE hardware and 13th generation Dell hardware running firmware version **2.75.75.75** or greater. Deploying with managed Secure Boot does not require Redfish virtual media. See "Configuring managed Secure Boot" in the "Setting up the environment for an OpenShift installation" section for details.



NOTE

Red Hat does not support Secure Boot with self-generated keys.

16.2.2. Planning a bare metal cluster for OpenShift Virtualization

If you will use OpenShift Virtualization, it is important to be aware of several requirements before you install your bare metal cluster.

- If you want to use live migration features, you must have multiple worker nodes *at the time of*

cluster installation. This is because live migration requires the cluster-level high availability (HA) flag to be set to true. The HA flag is set when a cluster is installed and cannot be changed afterwards. If there are fewer than two worker nodes defined when you install your cluster, the HA flag is set to false for the life of the cluster.



NOTE

You can install OpenShift Virtualization on a single-node cluster, but single-node OpenShift does not support high availability.

- Live migration requires shared storage. Storage for OpenShift Virtualization must support and use the ReadWriteMany (RWX) access mode.
- If you plan to use Single Root I/O Virtualization (SR-IOV), ensure that your network interface controllers (NICs) are supported by OpenShift Container Platform.

Additional resources

- [Preparing your cluster for OpenShift Virtualization](#)
- [About Single Root I/O Virtualization \(SR-IOV\) hardware networks](#)
- [Connecting a virtual machine to an SR-IOV network](#)

16.2.3. Firmware requirements for installing with virtual media

The installation program for installer-provisioned OpenShift Container Platform clusters validates the hardware and firmware compatibility with Redfish virtual media. The installation program does not begin installation on a node if the node firmware is not compatible. The following tables list the minimum firmware versions tested and verified to work for installer-provisioned OpenShift Container Platform clusters deployed by using Redfish virtual media.



NOTE

Red Hat does not test every combination of firmware, hardware, or other third-party components. For further information about third-party support, see [Red Hat third-party support policy](#). For information about updating the firmware, see the hardware documentation for the nodes or contact the hardware vendor.

Table 16.1. Firmware compatibility for HP hardware with Redfish virtual media

Model	Management	Firmware versions
10th Generation	iLO5	2.63 or later

Table 16.2. Firmware compatibility for Dell hardware with Redfish virtual media

Model	Management	Firmware versions
15th Generation	iDRAC 9	v6.10.30.00 and v7.00.60.00

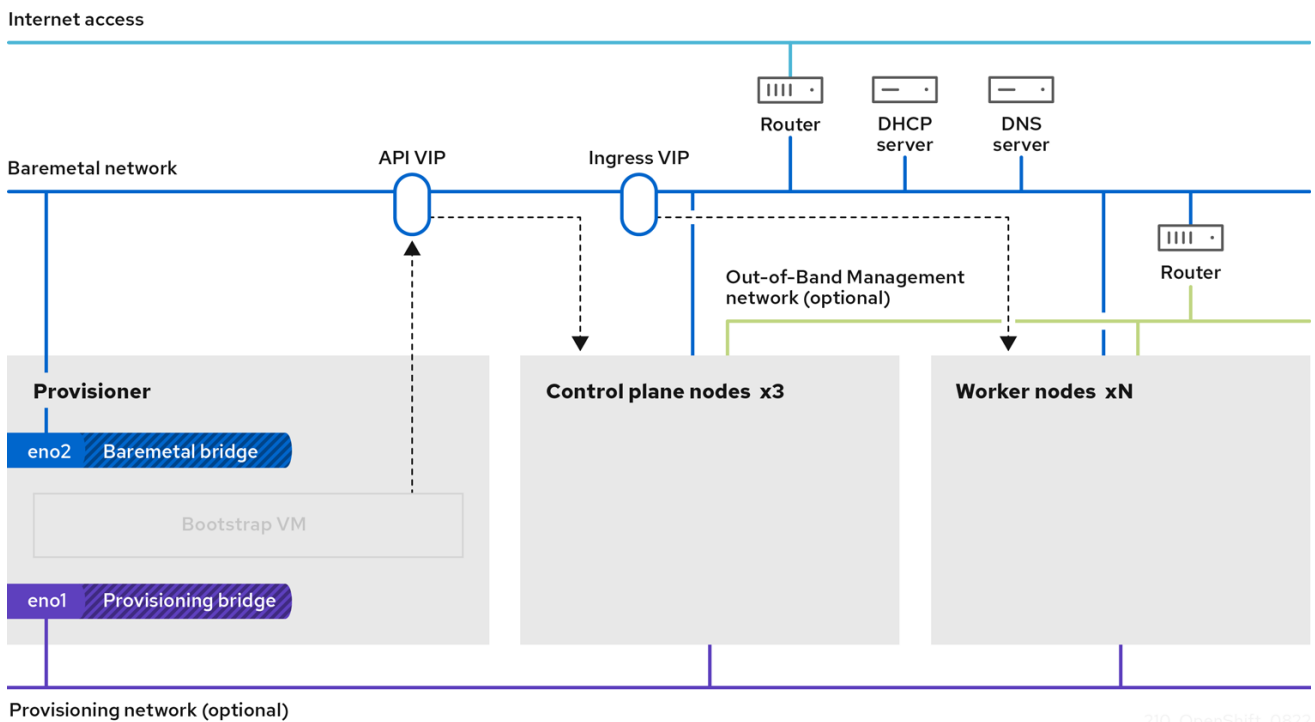
Model	Management	Firmware versions
14th Generation	iDRAC 9	v6.10.30.00
13th Generation	iDRAC 8	v2.75.75.75 or later

Additional resources

[Unable to discover new bare metal hosts using the BMC](#)

16.2.4. Network requirements

Installer-provisioned installation of OpenShift Container Platform involves several network requirements. First, installer-provisioned installation involves an optional non-routable **provisioning** network for provisioning the operating system on each bare metal node. Second, installer-provisioned installation involves a routable **baremetal** network.



16.2.4.1. Ensuring required ports are open

Certain ports must be open between cluster nodes for installer-provisioned installations to complete successfully. In certain situations, such as using separate subnets for far edge worker nodes, you must ensure that the nodes in these subnets can communicate with nodes in the other subnets on the following required ports.

Table 16.3. Required ports

Port	Description
67,68	When using a provisioning network, cluster nodes access the dnsmasq DHCP server over their provisioning network interfaces using ports 67 and 68 .
69	When using a provisioning network, cluster nodes communicate with the TFTP server on port 69 using their provisioning network interfaces. The TFTP server runs on the bootstrap VM. The bootstrap VM runs on the provisioner node.
80	When not using the image caching option or when using virtual media, the provisioner node must have port 80 open on the baremetal machine network interface to stream the Red Hat Enterprise Linux CoreOS (RHCOS) image from the provisioner node to the cluster nodes.
123	The cluster nodes must access the NTP server on port 123 using the baremetal machine network.
5050	The Ironic Inspector API runs on the control plane nodes and listens on port 5050 . The Inspector API is responsible for hardware introspection, which collects information about the hardware characteristics of the bare metal nodes.
5051	Port 5050 uses port 5051 as a proxy.
6180	When deploying with virtual media and not using TLS, the provisioner node and the control plane nodes must have port 6180 open on the baremetal machine network interface so that the baseboard management controller (BMC) of the worker nodes can access the RHCOS image. Starting with OpenShift Container Platform 4.13, the default HTTP port is 6180 .
6183	When deploying with virtual media and using TLS, the provisioner node and the control plane nodes must have port 6183 open on the baremetal machine network interface so that the BMC of the worker nodes can access the RHCOS image.

Port	Description
6385	The Ironic API server runs initially on the bootstrap VM and later on the control plane nodes and listens on port 6385 . The Ironic API allows clients to interact with Ironic for bare metal node provisioning and management, including operations like enrolling new nodes, managing their power state, deploying images, and cleaning the hardware.
6388	Port 6385 uses port 6388 as a proxy.
8080	When using image caching without TLS, port 8080 must be open on the provisioner node and accessible by the BMC interfaces of the cluster nodes.
8083	When using the image caching option with TLS, port 8083 must be open on the provisioner node and accessible by the BMC interfaces of the cluster nodes.
9999	By default, the Ironic Python Agent (IPA) listens on TCP port 9999 for API calls from the Ironic conductor service. This port is used for communication between the bare metal node where IPA is running and the Ironic conductor service.

16.2.4.2. Increase the network MTU

Before deploying OpenShift Container Platform, increase the network maximum transmission unit (MTU) to 1500 or more. If the MTU is lower than 1500, the Ironic image that is used to boot the node might fail to communicate with the Ironic inspector pod, and inspection will fail. If this occurs, installation stops because the nodes are not available for installation.

16.2.4.3. Configuring NICs

OpenShift Container Platform deploys with two networks:

- **provisioning**: The **provisioning** network is an optional non-routable network used for provisioning the underlying operating system on each node that is a part of the OpenShift Container Platform cluster. The network interface for the **provisioning** network on each cluster node must have the BIOS or UEFI configured to PXE boot.

The **provisioningNetworkInterface** configuration setting specifies the **provisioning** network NIC name on the control plane nodes, which must be identical on the control plane nodes. The **bootMACAddress** configuration setting provides a means to specify a particular NIC on each node for the **provisioning** network.

The **provisioning** network is optional, but it is required for PXE booting. If you deploy without a **provisioning** network, you must use a virtual media BMC addressing option such as **redfish-virtualmedia** or **idrac-virtualmedia**.

- **baremetal**: The **baremetal** network is a routable network. You can use any NIC to interface with the **baremetal** network provided the NIC is not configured to use the **provisioning** network.



IMPORTANT

When using a VLAN, each NIC must be on a separate VLAN corresponding to the appropriate network.

16.2.4.4. DNS requirements

Clients access the OpenShift Container Platform cluster nodes over the **baremetal** network. A network administrator must configure a subdomain or subzone where the canonical name extension is the cluster name.

```
<cluster_name>.<base_domain>
```

For example:

```
test-cluster.example.com
```

OpenShift Container Platform includes functionality that uses cluster membership information to generate A/AAAA records. This resolves the node names to their IP addresses. After the nodes are registered with the API, the cluster can disperse node information without using CoreDNS-mDNS. This eliminates the network traffic associated with multicast DNS.

In OpenShift Container Platform deployments, DNS name resolution is required for the following components:

- The Kubernetes API
- The OpenShift Container Platform application wildcard ingress API

A/AAAA records are used for name resolution and PTR records are used for reverse name resolution. Red Hat Enterprise Linux CoreOS (RHCOS) uses the reverse records or DHCP to set the hostnames for all the nodes.

Installer-provisioned installation includes functionality that uses cluster membership information to generate A/AAAA records. This resolves the node names to their IP addresses. In each record, **<cluster_name>** is the cluster name and **<base_domain>** is the base domain that you specify in the **install-config.yaml** file. A complete DNS record takes the form: **<component>.<cluster_name>.<base_domain>.**

Table 16.4. Required DNS records

Component	Record	Description
Kubernetes API	api.<cluster_name>.<base_domain>.	An A/AAAA record and a PTR record identify the API load balancer. These records must be resolvable by both clients external to the cluster and from all the nodes within the cluster.

Component	Record	Description
Routes	*.apps.<cluster_name>.<base_domain>	<p>The wildcard A/AAAA record refers to the application ingress load balancer. The application ingress load balancer targets the nodes that run the Ingress Controller pods. The Ingress Controller pods run on the worker nodes by default. These records must be resolvable by both clients external to the cluster and from all the nodes within the cluster.</p> <p>For example, console-openshift-console.apps.<cluster_name>.<base_domain> is used as a wildcard route to the OpenShift Container Platform console.</p>

TIP

You can use the **dig** command to verify DNS resolution.

16.2.4.5. Dynamic Host Configuration Protocol (DHCP) requirements

By default, installer-provisioned installation deploys **ironic-dnsmasq** with DHCP enabled for the **provisioning** network. No other DHCP servers should be running on the **provisioning** network when the **provisioningNetwork** configuration setting is set to **managed**, which is the default value. If you have a DHCP server running on the **provisioning** network, you must set the **provisioningNetwork** configuration setting to **unmanaged** in the **install-config.yaml** file.

Network administrators must reserve IP addresses for each node in the OpenShift Container Platform cluster for the **baremetal** network on an external DHCP server.

16.2.4.6. Reserving IP addresses for nodes with the DHCP server

For the **baremetal** network, a network administrator must reserve a number of IP addresses, including:

1. Two unique virtual IP addresses.
 - One virtual IP address for the API endpoint.
 - One virtual IP address for the wildcard ingress endpoint.
2. One IP address for the provisioner node.
3. One IP address for each control plane node.
4. One IP address for each worker node, if applicable.

**RESERVING IP ADDRESSES SO THEY BECOME STATIC IP ADDRESSES**

Some administrators prefer to use static IP addresses so that each node's IP address remains constant in the absence of a DHCP server. To configure static IP addresses with NMState, see "(Optional) Configuring node network interfaces" in the "Setting up the environment for an OpenShift installation" section.



NETWORKING BETWEEN EXTERNAL LOAD BALANCERS AND CONTROL PLANE NODES

External load balancing services and the control plane nodes must run on the same L2 network, and on the same VLAN when using VLANs to route traffic between the load balancing services and the control plane nodes.



IMPORTANT

The storage interface requires a DHCP reservation or a static IP.

The following table provides an exemplary embodiment of fully qualified domain names. The API and Nameserver addresses begin with canonical name extensions. The hostnames of the control plane and worker nodes are exemplary, so you can use any host naming convention you prefer.

Usage	Host Name	IP
API	api.<cluster_name>.<base_domain>	<ip>
Ingress LB (apps)	*.apps.<cluster_name>.<base_domain>	<ip>
Provisioner node	provisioner.<cluster_name>.<base_domain>	<ip>
Control-plane-0	openshift-control-plane-0.<cluster_name>.<base_domain>	<ip>
Control-plane-1	openshift-control-plane-1.<cluster_name>.<base_domain>	<ip>
Control-plane-2	openshift-control-plane-2.<cluster_name>.<base_domain>	<ip>
Worker-0	openshift-worker-0.<cluster_name>.<base_domain>	<ip>
Worker-1	openshift-worker-1.<cluster_name>.<base_domain>	<ip>
Worker-n	openshift-worker-n.<cluster_name>.<base_domain>	<ip>



NOTE

If you do not create DHCP reservations, the installer requires reverse DNS resolution to set the hostnames for the Kubernetes API node, the provisioner node, the control plane nodes, and the worker nodes.

16.2.4.7. Provisioner node requirements


You must specify the MAC address for the provisioner node in your installation configuration. The

bootMacAddress specification is typically associated with PXE network booting. However, the Ironic provisioning service also requires the **bootMacAddress** specification to identify nodes during the inspection of the cluster, or during node redeployment in the cluster.

The provisioner node requires layer 2 connectivity for network booting, DHCP and DNS resolution, and local network communication. The provisioner node requires layer 3 connectivity for virtual media booting.

16.2.4.8. Network Time Protocol (NTP)

Each OpenShift Container Platform node in the cluster must have access to an NTP server. OpenShift Container Platform nodes use NTP to synchronize their clocks. For example, cluster nodes use SSL certificates that require validation, which might fail if the date and time between the nodes are not in sync.



IMPORTANT

Define a consistent clock date and time format in each cluster node's BIOS settings, or installation might fail.

You can reconfigure the control plane nodes to act as NTP servers on disconnected clusters, and reconfigure worker nodes to retrieve time from the control plane nodes.

16.2.4.9. Port access for the out-of-band management IP address

The out-of-band management IP address is on a separate network from the node. To ensure that the out-of-band management can communicate with the provisioner node during installation, the out-of-band management IP address must be granted access to port **6180** on the provisioner node and on the OpenShift Container Platform control plane nodes. TLS port **6183** is required for virtual media installation, for example, by using Redfish.

Additional resources

- [Using DNS forwarding](#)

16.2.5. Configuring nodes

Configuring nodes when using the provisioning network

Each node in the cluster requires the following configuration for proper installation.



WARNING

A mismatch between nodes will cause an installation failure.

While the cluster nodes can contain more than two NICs, the installation process only focuses on the first two NICs. In the following table, NIC1 is a non-routable network (**provisioning**) that is only used for the installation of the OpenShift Container Platform cluster.

NIC	Network	VLAN
NIC1	provisioning	<provisioning_vlan>
NIC2	baremetal	<baremetal_vlan>

The Red Hat Enterprise Linux (RHEL) 9.x installation process on the provisioner node might vary. To install Red Hat Enterprise Linux (RHEL) 9.x using a local Satellite server or a PXE server, PXE-enable NIC2.

PXE	Boot order
NIC1 PXE-enabled provisioning network	1
NIC2 baremetal network. PXE-enabled is optional.	2

**NOTE**

Ensure PXE is disabled on all other NICs.

Configure the control plane and worker nodes as follows:

PXE	Boot order
NIC1 PXE-enabled (provisioning network)	1

Configuring nodes without the provisioning network

The installation process requires one NIC:

NIC	Network	VLAN
NICx	baremetal	<baremetal_vlan>


NICx is a routable network (**baremetal**) that is used for the installation of the OpenShift Container Platform cluster, and routable to the internet.

**IMPORTANT**

The **provisioning** network is optional, but it is required for PXE booting. If you deploy without a **provisioning** network, you must use a virtual media BMC addressing option such as **redfish-virtualmedia** or **idrac-virtualmedia**.

Configuring nodes for Secure Boot manually

Secure Boot prevents a node from booting unless it verifies the node is using only trusted software, such as UEFI firmware drivers, EFI applications, and the operating system.


**NOTE**

Red Hat only supports manually configured Secure Boot when deploying with Redfish virtual media.

To enable Secure Boot manually, refer to the hardware guide for the node and execute the following:

Procedure

1. Boot the node and enter the BIOS menu.
2. Set the node's boot mode to **UEFI Enabled**.
3. Enable Secure Boot.

**IMPORTANT**


Red Hat does not support Secure Boot with self-generated keys.

16.2.6. Out-of-band management

Nodes typically have an additional NIC used by the baseboard management controllers (BMCs). These BMCs must be accessible from the provisioner node.

Each node must be accessible via out-of-band management. When using an out-of-band management network, the provisioner node requires access to the out-of-band management network for a successful OpenShift Container Platform installation.

The out-of-band management setup is out of scope for this document. Using a separate management network for out-of-band management can enhance performance and improve security. However, using the provisioning network or the bare metal network are valid options.

**NOTE**

The bootstrap VM features a maximum of two network interfaces. If you configure a separate management network for out-of-band management, and you are using a provisioning network, the bootstrap VM requires routing access to the management network through one of the network interfaces. In this scenario, the bootstrap VM can then access three networks:

- the bare metal network
- the provisioning network
- the management network routed through one of the network interfaces

16.2.7. Required data for installation

Prior to the installation of the OpenShift Container Platform cluster, gather the following information from all cluster nodes:

- Out-of-band management IP
 - Examples

- Dell (iDRAC) IP
- HP (iLO) IP
- Fujitsu (iRMC) IP

When using the provisioning network

- NIC (**provisioning**) MAC address
- NIC (**baremetal**) MAC address

When omitting the provisioning network

- NIC (**baremetal**) MAC address

16.2.8. Validation checklist for nodes

When using the provisioning network

- NIC1 VLAN is configured for the **provisioning** network.
- NIC1 for the **provisioning** network is PXE-enabled on the provisioner, control plane, and worker nodes.
- NIC2 VLAN is configured for the **baremetal** network.
- PXE has been disabled on all other NICs.
- DNS is configured with API and Ingress endpoints.
- Control plane and worker nodes are configured.
- All nodes accessible via out-of-band management.
- (Optional) A separate management network has been created.
- Required data for installation.

When omitting the provisioning network

- NIC1 VLAN is configured for the **baremetal** network.
- DNS is configured with API and Ingress endpoints.
- Control plane and worker nodes are configured.
- All nodes accessible via out-of-band management.
- (Optional) A separate management network has been created.
- Required data for installation.

16.3. SETTING UP THE ENVIRONMENT FOR AN OPENSIFT INSTALLATION

16.3.1. Installing RHEL on the provisioner node

With the configuration of the prerequisites complete, the next step is to install RHEL 9.x on the provisioner node. The installer uses the provisioner node as the orchestrator while installing the OpenShift Container Platform cluster. For the purposes of this document, installing RHEL on the provisioner node is out of scope. However, options include but are not limited to using a RHEL Satellite server, PXE, or installation media.

16.3.2. Preparing the provisioner node for OpenShift Container Platform installation

Perform the following steps to prepare the environment.

Procedure

1. Log in to the provisioner node via **ssh**.
2. Create a non-root user (**kni**) and provide that user with **sudo** privileges:

```
# useradd kni
```

```
# passwd kni
```

```
# echo "kni ALL=(root) NOPASSWD:ALL" | tee -a /etc/sudoers.d/kni
```

```
# chmod 0440 /etc/sudoers.d/kni
```

3. Create an **ssh** key for the new user:

```
# su - kni -c "ssh-keygen -t ed25519 -f /home/kni/.ssh/id_rsa -N ""
```

4. Log in as the new user on the provisioner node:

```
# su - kni
```

5. Use Red Hat Subscription Manager to register the provisioner node:

```
$ sudo subscription-manager register --username=<user> --password=<pass> --auto-attach
```

```
$ sudo subscription-manager repos --enable=rhel-9-for-<architecture>-appstream-rpms --enable=rhel-9-for-<architecture>-baseos-rpms
```



NOTE

For more information about Red Hat Subscription Manager, see [Using and Configuring Red Hat Subscription Manager](#).

6. Install the following packages:

```
$ sudo dnf install -y libvirt qemu-kvm mkisofs python3-devel jq ipmitool
```

7. Modify the user to add the **libvirt** group to the newly created user:

```
$ sudo usermod --append --groups libvirt <user>
```

- Restart **firewalld** and enable the **http** service:

```
$ sudo systemctl start firewalld
```

```
$ sudo firewall-cmd --zone=public --add-service=http --permanent
```

```
$ sudo firewall-cmd --reload
```

- Start and enable the **libvirtd** service:

```
$ sudo systemctl enable libvirtd --now
```

- Create the **default** storage pool and start it:

```
$ sudo virsh pool-define-as --name default --type dir --target /var/lib/libvirt/images
```

```
$ sudo virsh pool-start default
```

```
$ sudo virsh pool-autostart default
```

- Create a **pull-secret.txt** file:

```
$ vim pull-secret.txt
```

In a web browser, navigate to [Install OpenShift on Bare Metal with installer-provisioned infrastructure](#). Click **Copy pull secret**. Paste the contents into the **pull-secret.txt** file and save the contents in the **kni** user's home directory.

16.3.3. Checking NTP server synchronization

The OpenShift Container Platform installation program installs the **chrony** Network Time Protocol (NTP) service on the cluster nodes. To complete installation, each node must have access to an NTP time server. You can verify NTP server synchronization by using the **chrony** service.

For disconnected clusters, you must configure the NTP servers on the control plane nodes. For more information see the *Additional resources* section.

Prerequisites

- You installed the **chrony** package on the target node.

Procedure

- Log in to the node by using the **ssh** command.
- View the NTP servers available to the node by running the following command:

```
$ chronyc sources
```

Example output

```

MS Name/IP address      Stratum Poll Reach LastRx Last sample
=====
=====
^+ time.cloudflare.com   3 10 377 187 -209us[-209us] +/- 32ms
^+ t1.time.ir2.yahoo.com 2 10 377 185 -4382us[-4382us] +/- 23ms
^+ time.cloudflare.com   3 10 377 198 -996us[-1220us] +/- 33ms
^* brenbox.westnet.ie    1 10 377 193 -9538us[-9761us] +/- 24ms

```

- Use the **ping** command to ensure that the node can access an NTP server, for example:

```
$ ping time.cloudflare.com
```

Example output

```

PING time.cloudflare.com (162.159.200.123) 56(84) bytes of data.
64 bytes from time.cloudflare.com (162.159.200.123): icmp_seq=1 ttl=54 time=32.3 ms
64 bytes from time.cloudflare.com (162.159.200.123): icmp_seq=2 ttl=54 time=30.9 ms
64 bytes from time.cloudflare.com (162.159.200.123): icmp_seq=3 ttl=54 time=36.7 ms
...

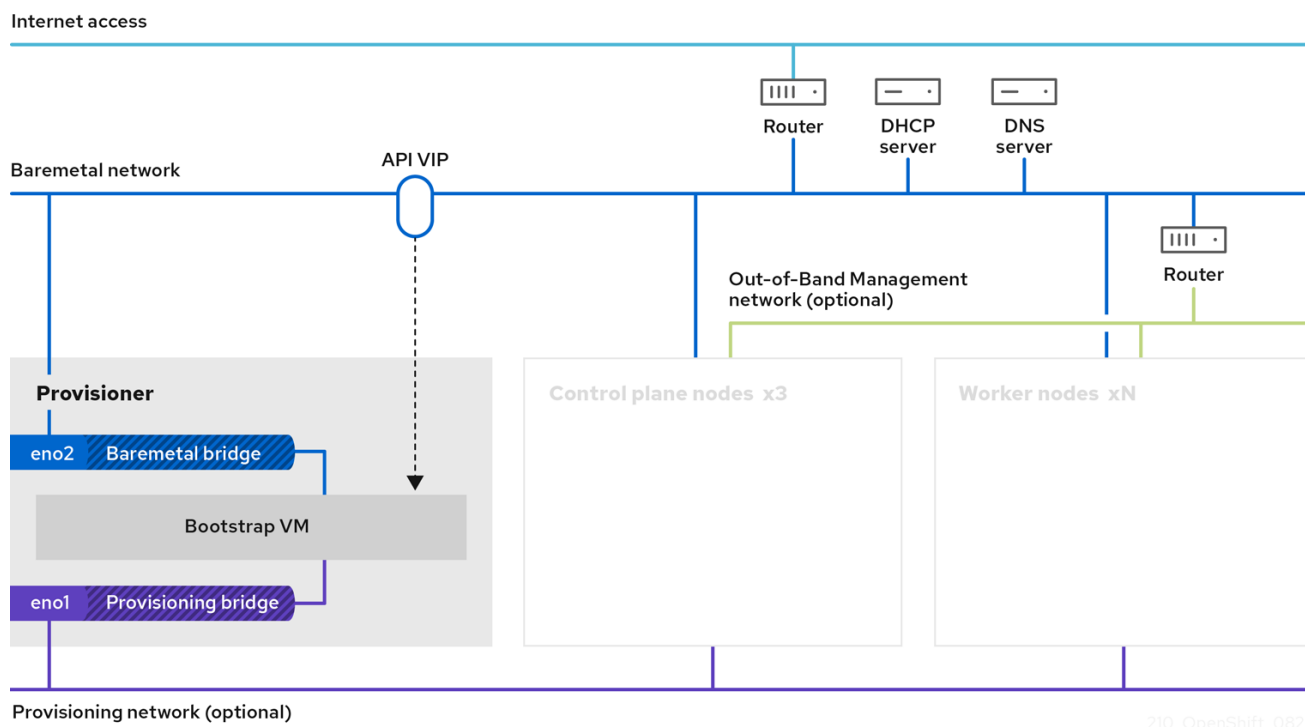
```

Additional resources

- [Optional: Configuring NTP for disconnected clusters](#)
- [Network Time Protocol \(NTP\)](#)

16.3.4. Configuring networking

Before installation, you must configure the networking on the provisioner node. Installer-provisioned clusters deploy with a bare-metal bridge and network, and an optional provisioning bridge and network.



210_OpenShift_0822

**NOTE**

You can also configure networking from the web console.

Procedure

1. Export the bare-metal network NIC name:

```
$ export PUB_CONN=<baremetal_nic_name>
```

2. Configure the bare-metal network:

**NOTE**

The SSH connection might disconnect after executing these steps.

```
$ sudo nohup bash -c "
  nmcli con down \"$PUB_CONN\"
  nmcli con delete \"$PUB_CONN\"
  # RHEL 8.1 appends the word \"System\" in front of the connection, delete in case it exists
  nmcli con down \"System $PUB_CONN\"
  nmcli con delete \"System $PUB_CONN\"
  nmcli connection add ifname baremetal type bridge con-name baremetal bridge.stp no
  nmcli con add type bridge-slave ifname \"$PUB_CONN\" master baremetal
  pkill dhclient;dhclient baremetal
"
```

3. Optional: If you are deploying with a provisioning network, export the provisioning network NIC name:

```
$ export PROV_CONN=<prov_nic_name>
```

4. Optional: If you are deploying with a provisioning network, configure the provisioning network:

```
$ sudo nohup bash -c "
  nmcli con down \"\$PROV_CONN\"
  nmcli con delete \"\$PROV_CONN\"
  nmcli connection add ifname provisioning type bridge con-name provisioning
  nmcli con add type bridge-slave ifname \"\$PROV_CONN\" master provisioning
  nmcli connection modify provisioning ipv6.addresses fd00:1101::1/64 ipv6.method manual
  nmcli con down provisioning
  nmcli con up provisioning
"
```



NOTE

The ssh connection might disconnect after executing these steps.

The IPv6 address can be any address as long as it is not routable via the bare-metal network.

Ensure that UEFI is enabled and UEFI PXE settings are set to the IPv6 protocol when using IPv6 addressing.

5. Optional: If you are deploying with a provisioning network, configure the IPv4 address on the provisioning network connection:

```
$ nmcli connection modify provisioning ipv4.addresses 172.22.0.254/24 ipv4.method manual
```

6. **ssh** back into the **provisioner** node (if required):

```
# ssh kni@provisioner.<cluster-name>.<domain>
```

7. Verify the connection bridges have been properly created:

```
$ sudo nmcli con show
```

NAME	UUID	TYPE	DEVICE
baremetal	4d5133a5-8351-4bb9-bfd4-3af264801530	bridge	baremetal
provisioning	43942805-017f-4d7d-a2c2-7cb3324482ed	bridge	provisioning
virbr0	d9bca40f-eee1-410b-8879-a2d4bb0465e7	bridge	virbr0
bridge-slave-eno1	76a8ed50-c7e5-4999-b4f6-6d9014dd0812	ethernet	eno1
bridge-slave-eno2	f31c3353-54b7-48de-893a-02d2b34c4736	ethernet	eno2

16.3.5. Establishing communication between subnets

In a typical OpenShift Container Platform cluster setup, all nodes, including the control plane and compute nodes, reside in the same network. However, for edge computing scenarios, it can be beneficial to locate compute nodes closer to the edge. This often involves using different network segments or subnets for the remote nodes than the subnet used by the control plane and local compute nodes. Such a setup can reduce latency for the edge and allow for enhanced scalability.

Before installing OpenShift Container Platform, you must configure the network properly to ensure that the edge subnets containing the remote nodes can reach the subnet containing the control plane nodes and receive traffic from the control plane too.

You can run control plane nodes in the same subnet or multiple subnets by configuring a user-managed load balancer in place of the default load balancer. With a multiple subnet environment, you can reduce the risk of your OpenShift Container Platform cluster from failing because of a hardware failure or a network outage. For more information, see "Services for a user-managed load balancer" and "Configuring a user-managed load balancer".

Running control plane nodes in a multiple subnet environment requires completion of the following key tasks:

- Configuring a user-managed load balancer instead of the default load balancer by specifying **UserManaged** in the **loadBalancer.type** parameter of the **install-config.yaml** file.
- Configuring a user-managed load balancer address in the **ingressVIPs** and **apiVIPs** parameters of the **install-config.yaml** file.
- Adding the multiple subnet Classless Inter-Domain Routing (CIDR) and the user-managed load balancer IP addresses to the **networking.machineNetworks** parameter in the **install-config.yaml** file.



NOTE

Deploying a cluster with multiple subnets requires using virtual media, such as **redfish-virtualmedia** and **idrac-virtualmedia**.

This procedure details the network configuration required to allow the remote compute nodes in the second subnet to communicate effectively with the control plane nodes in the first subnet and to allow the control plane nodes in the first subnet to communicate effectively with the remote compute nodes in the second subnet.

In this procedure, the cluster spans two subnets:

- The first subnet (**10.0.0.0**) contains the control plane and local compute nodes.
- The second subnet (**192.168.0.0**) contains the edge compute nodes.

Procedure

1. Configure the first subnet to communicate with the second subnet:

- a. Log in as **root** to a control plane node by running the following command:

```
$ sudo su -
```

- b. Get the name of the network interface by running the following command:

```
# nmcli dev status
```

- c. Add a route to the second subnet (**192.168.0.0**) via the gateway by running the following command:

```
# nmcli connection modify <interface_name> +ipv4.routes "192.168.0.0/24 via <gateway>"
```

Replace **<interface_name>** with the interface name. Replace **<gateway>** with the IP address of the actual gateway.

Example

```
# nmcli connection modify eth0 +ipv4.routes "192.168.0.0/24 via 192.168.0.1"
```

- d. Apply the changes by running the following command:

```
# nmcli connection up <interface_name>
```

Replace **<interface_name>** with the interface name.

- e. Verify the routing table to ensure the route has been added successfully:

```
# ip route
```

- f. Repeat the previous steps for each control plane node in the first subnet.



NOTE

Adjust the commands to match your actual interface names and gateway.

2. Configure the second subnet to communicate with the first subnet:

- a. Log in as **root** to a remote compute node by running the following command:

```
$ sudo su -
```

- b. Get the name of the network interface by running the following command:

```
# nmcli dev status
```

- c. Add a route to the first subnet (**10.0.0.0**) via the gateway by running the following command:

```
# nmcli connection modify <interface_name> +ipv4.routes "10.0.0.0/24 via <gateway>"
```

Replace **<interface_name>** with the interface name. Replace **<gateway>** with the IP address of the actual gateway.

Example

```
# nmcli connection modify eth0 +ipv4.routes "10.0.0.0/24 via 10.0.0.1"
```

- d. Apply the changes by running the following command:


```
# nmcli connection up <interface_name>
```

Replace **<interface_name>** with the interface name.

- e. Verify the routing table to ensure the route has been added successfully by running the following command:

```
# ip route
```

- f. Repeat the previous steps for each compute node in the second subnet.



NOTE

Adjust the commands to match your actual interface names and gateway.

3. After you have configured the networks, test the connectivity to ensure the remote nodes can reach the control plane nodes and the control plane nodes can reach the remote nodes.

- a. From the control plane nodes in the first subnet, ping a remote node in the second subnet by running the following command:

```
$ ping <remote_node_ip_address>
```

If the ping is successful, it means the control plane nodes in the first subnet can reach the remote nodes in the second subnet. If you do not receive a response, review the network configurations and repeat the procedure for the node.

- b. From the remote nodes in the second subnet, ping a control plane node in the first subnet by running the following command:

```
$ ping <control_plane_node_ip_address>
```

If the ping is successful, it means the remote compute nodes in the second subnet can reach the control plane in the first subnet. If you do not receive a response, review the network configurations and repeat the procedure for the node.

16.3.6. Retrieving the OpenShift Container Platform installer

Use the **stable-4.x** version of the installation program and your selected architecture to deploy the generally available stable version of OpenShift Container Platform:

```
$ export VERSION=stable-4.16
```

```
$ export RELEASE_ARCH=<architecture>
```

```
$ export RELEASE_IMAGE=$(curl -s https://mirror.openshift.com/pub/openshift-
v4/$RELEASE_ARCH/clients/ocp/$VERSION/release.txt | grep 'Pull From: quay.io' | awk -F ' ' '{print
$3}')
```

16.3.7. Extracting the OpenShift Container Platform installer

After retrieving the installer, the next step is to extract it.

Procedure

1. Set the environment variables:

```
$ export cmd=openshift-baremetal-install
```

```
$ export pullsecret_file=~/.pull-secret.txt
```

```
$ export extract_dir=$(pwd)
```

2. Get the **oc** binary:

```
$ curl -s https://mirror.openshift.com/pub/openshift-v4/clients/ocp/$VERSION/openshift-client-linux.tar.gz | tar zxvf - oc
```

3. Extract the installer:

```
$ sudo cp oc /usr/local/bin
```

```
$ oc adm release extract --registry-config "${pullsecret_file}" --command=$cmd --to "${extract_dir}" ${RELEASE_IMAGE}
```

```
$ sudo cp openshift-baremetal-install /usr/local/bin
```

16.3.8. Optional: Creating an RHCOS images cache

To employ image caching, you must download the Red Hat Enterprise Linux CoreOS (RHCOS) image used by the bootstrap VM to provision the cluster nodes. Image caching is optional, but it is especially useful when running the installation program on a network with limited bandwidth.



NOTE

The installation program no longer needs the **clusterOSImage** RHCOS image because the correct image is in the release payload.

If you are running the installation program on a network with limited bandwidth and the RHCOS images download takes more than 15 to 20 minutes, the installation program will timeout. Caching images on a web server will help in such scenarios.



WARNING

If you enable TLS for the HTTPD server, you must confirm the root certificate is signed by an authority trusted by the client and verify the trusted certificate chain between your OpenShift Container Platform hub and spoke clusters and the HTTPD server. Using a server configured with an untrusted certificate prevents the images from being downloaded to the image creation service. Using untrusted HTTPS servers is not supported.

Install a container that contains the images.

Procedure

1. Install **podman**:

```
$ sudo dnf install -y podman
```

2. Open firewall port **8080** to be used for RHCOS image caching:

```
$ sudo firewall-cmd --add-port=8080/tcp --zone=public --permanent
```

```
$ sudo firewall-cmd --reload
```

3. Create a directory to store the **bootstraposimage**:

```
$ mkdir /home/kni/rhcos_image_cache
```

4. Set the appropriate SELinux context for the newly created directory:

```
$ sudo semanage fcontext -a -t httpd_sys_content_t "/home/kni/rhcos_image_cache(/.*)?"
```

```
$ sudo restorecon -Rv /home/kni/rhcos_image_cache/
```

5. Get the URI for the RHCOS image that the installation program will deploy on the bootstrap VM:

```
$ export RHCOS_QEMU_URI=$(/usr/local/bin/openshift-baremetal-install coreos print-stream-json | jq -r --arg ARCH "$(arch)"
'.architectures[$ARCH].artifacts.qemu.formats["qcow2.gz"].disk.location')
```

6. Get the name of the image that the installation program will deploy on the bootstrap VM:

```
$ export RHCOS_QEMU_NAME=${RHCOS_QEMU_URI##*/}
```

7. Get the SHA hash for the RHCOS image that will be deployed on the bootstrap VM:

```
$ export RHCOS_QEMU_UNCOMPRESSED_SHA256=$(/usr/local/bin/openshift-baremetal-install coreos print-stream-json | jq -r --arg ARCH "$(arch)"
'.architectures[$ARCH].artifacts.qemu.formats["qcow2.gz"].disk["uncompressed-sha256"]')
```

8. Download the image and place it in the **/home/kni/rhcos_image_cache** directory:

```
$ curl -L ${RHCOS_QEMU_URI} -o
/home/kni/rhcos_image_cache/${RHCOS_QEMU_NAME}
```

9. Confirm SELinux type is of **httpd_sys_content_t** for the new file:

```
$ ls -Z /home/kni/rhcos_image_cache
```

10. Create the pod:

```
$ podman run -d --name rhcos_image_cache \ 1
-v /home/kni/rhcos_image_cache:/var/www/html \
-p 8080:8080/tcp \
registry.access.redhat.com/ubi9/httpd-24
```

- 1 Creates a caching webserver with the name **rhcos_image_cache**. This pod serves the **bootstrapOSImage** image in the **install-config.yaml** file for deployment.

11. Generate the **bootstrapOSImage** configuration:

```
$ export BAREMETAL_IP=$(ip addr show dev baremetal | awk '/inet /{print $2}' | cut -d"/" -f1)
```

```
$ export
BOOTSTRAP_OS_IMAGE="http://${BAREMETAL_IP}:8080/${RHCOS_QEMU_NAME}?
sha256=${RHCOS_QEMU_UNCOMPRESSED_SHA256}"
```

```
$ echo " bootstrapOSImage=${BOOTSTRAP_OS_IMAGE}"
```

12. Add the required configuration to the **install-config.yaml** file under **platform.baremetal**:


```
platform:
  baremetal:
    bootstrapOSImage: <bootstrap_os_image> 1
```

- 1 Replace **<bootstrap_os_image>** with the value of **\$BOOTSTRAP_OS_IMAGE**.

See the "Configuring the install-config.yaml file" section for additional details.

16.3.9. Services for a user-managed load balancer

You can configure an OpenShift Container Platform cluster to use a user-managed load balancer in place of the default load balancer.



IMPORTANT

Configuring a user-managed load balancer depends on your vendor's load balancer.

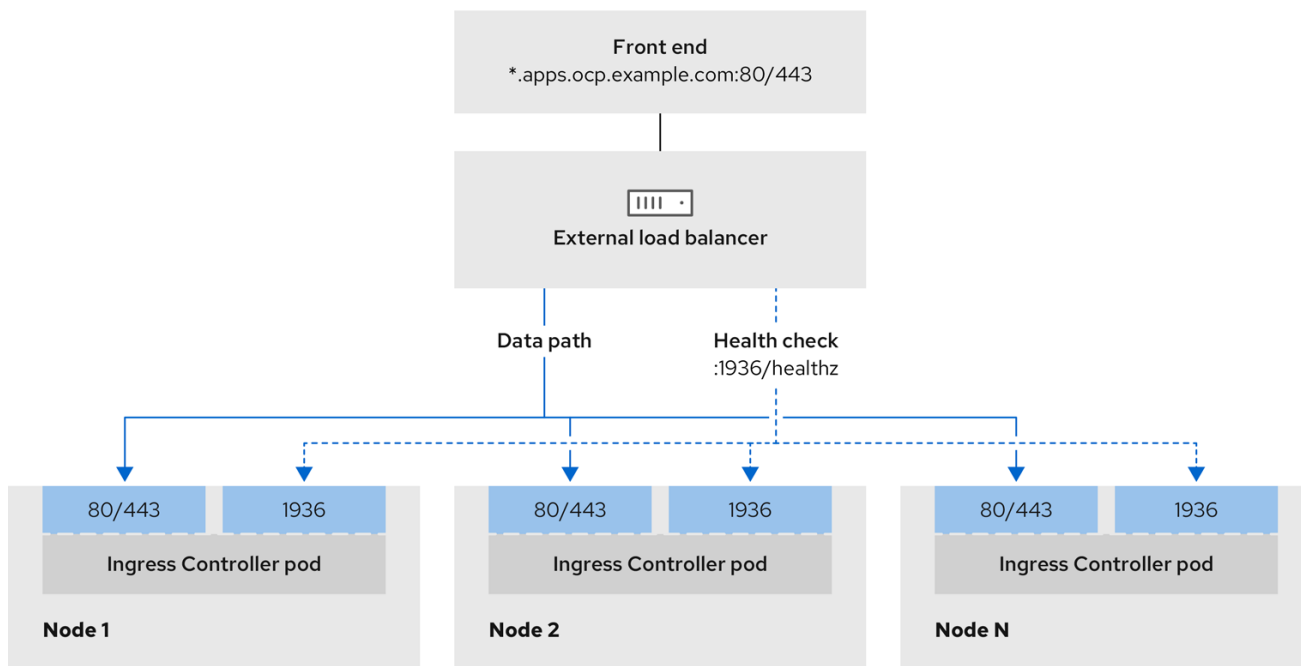
The information and examples in this section are for guideline purposes only. Consult the vendor documentation for more specific information about the vendor's load balancer.

Red Hat supports the following services for a user-managed load balancer:

- Ingress Controller
- OpenShift API
- OpenShift MachineConfig API

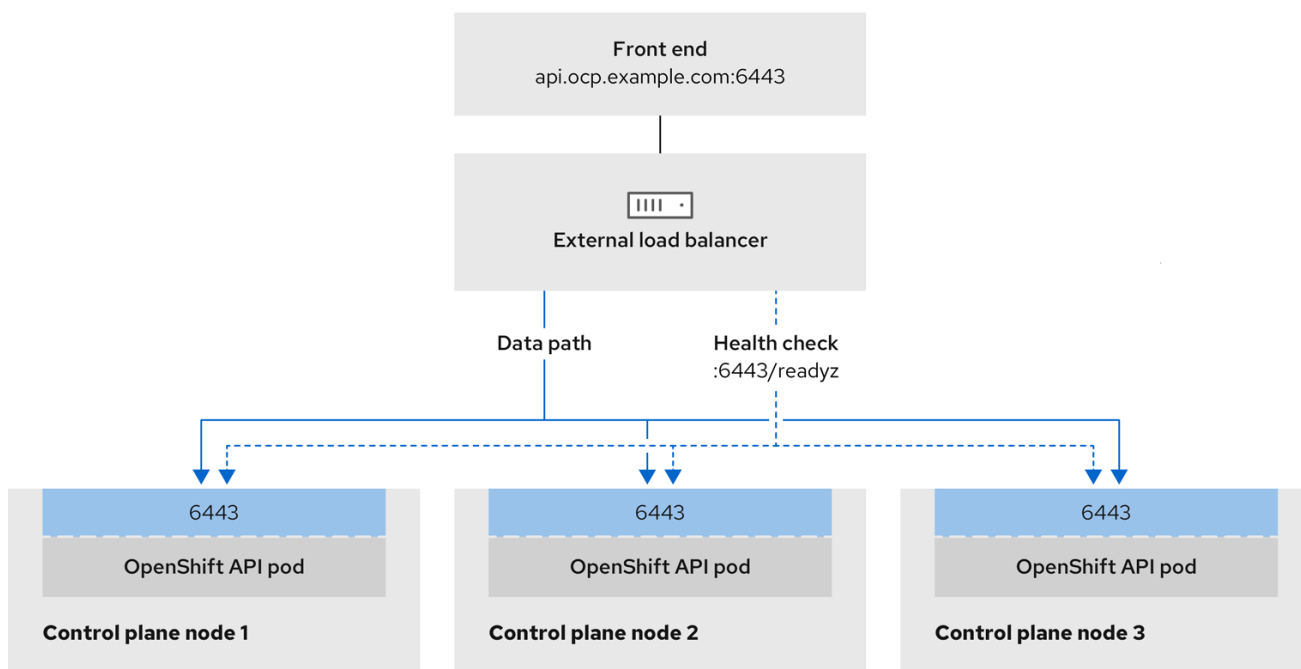
You can choose whether you want to configure one or all of these services for a user-managed load balancer. Configuring only the Ingress Controller service is a common configuration option. To better understand each service, view the following diagrams:

Figure 16.1. Example network workflow that shows an Ingress Controller operating in an OpenShift Container Platform environment




496_OpenShift_1223

Figure 16.2. Example network workflow that shows an OpenShift API operating in an OpenShift Container Platform environment



496_OpenShift_1223

Figure 16.3. Example network workflow that shows an OpenShift MachineConfig API operating in an OpenShift Container Platform environment



496_OpenShift_1223

The following configuration options are supported for user-managed load balancers:

- Use a node selector to map the Ingress Controller to a specific set of nodes. You must assign a static IP address to each node in this set, or configure each node to receive the same IP address from the Dynamic Host Configuration Protocol (DHCP). Infrastructure nodes commonly receive this type of configuration.
- Target all IP addresses on a subnet. This configuration can reduce maintenance overhead, because you can create and destroy nodes within those networks without reconfiguring the load balancer targets. If you deploy your ingress pods by using a machine set on a smaller network, such as a `/27` or `/28`, you can simplify your load balancer targets.

TIP

You can list all IP addresses that exist in a network by checking the machine config pool's resources.

Before you configure a user-managed load balancer for your OpenShift Container Platform cluster, consider the following information:

- For a front-end IP address, you can use the same IP address for the front-end IP address, the Ingress Controller's load balancer, and API load balancer. Check the vendor's documentation for this capability.
- For a back-end IP address, ensure that an IP address for an OpenShift Container Platform control plane node does not change during the lifetime of the user-managed load balancer. You can achieve this by completing one of the following actions:
 - Assign a static IP address to each control plane node.

- Configure each node to receive the same IP address from the DHCP every time the node requests a DHCP lease. Depending on the vendor, the DHCP lease might be in the form of an IP reservation or a static DHCP assignment.
- Manually define each node that runs the Ingress Controller in the user-managed load balancer for the Ingress Controller back-end service. For example, if the Ingress Controller moves to an undefined node, a connection outage can occur.

16.3.9.1. Configuring a user-managed load balancer

You can configure an OpenShift Container Platform cluster to use a user-managed load balancer in place of the default load balancer.



IMPORTANT

Before you configure a user-managed load balancer, ensure that you read the "Services for a user-managed load balancer" section.

Read the following prerequisites that apply to the service that you want to configure for your user-managed load balancer.



NOTE

MetaLB, which runs on a cluster, functions as a user-managed load balancer.

OpenShift API prerequisites

- You defined a front-end IP address.
- TCP ports 6443 and 22623 are exposed on the front-end IP address of your load balancer. Check the following items:
 - Port 6443 provides access to the OpenShift API service.
 - Port 22623 can provide ignition startup configurations to nodes.
- The front-end IP address and port 6443 are reachable by all users of your system with a location external to your OpenShift Container Platform cluster.
- The front-end IP address and port 22623 are reachable only by OpenShift Container Platform nodes.
- The load balancer backend can communicate with OpenShift Container Platform control plane nodes on port 6443 and 22623.

Ingress Controller prerequisites

- You defined a front-end IP address.
- TCP ports 443 and 80 are exposed on the front-end IP address of your load balancer.
- The front-end IP address, port 80 and port 443 are be reachable by all users of your system with a location external to your OpenShift Container Platform cluster.

- The front-end IP address, port 80 and port 443 are reachable to all nodes that operate in your OpenShift Container Platform cluster.
- The load balancer backend can communicate with OpenShift Container Platform nodes that run the Ingress Controller on ports 80, 443, and 1936.

Prerequisite for health check URL specifications

You can configure most load balancers by setting health check URLs that determine if a service is available or unavailable. OpenShift Container Platform provides these health checks for the OpenShift API, Machine Configuration API, and Ingress Controller backend services.

The following examples show health check specifications for the previously listed backend services:

Example of a Kubernetes API health check specification

```
Path: HTTPS:6443/readyz
Healthy threshold: 2
Unhealthy threshold: 2
Timeout: 10
Interval: 10
```

Example of a Machine Config API health check specification

```
Path: HTTPS:22623/healthz
Healthy threshold: 2
Unhealthy threshold: 2
Timeout: 10
Interval: 10
```

Example of an Ingress Controller health check specification

```
Path: HTTP:1936/healthz/ready
Healthy threshold: 2
Unhealthy threshold: 2
Timeout: 5
Interval: 10
```

Procedure

1. Configure the HAProxy Ingress Controller, so that you can enable access to the cluster from your load balancer on ports 6443, 22623, 443, and 80. Depending on your needs, you can specify the IP address of a single subnet or IP addresses from multiple subnets in your HAProxy configuration.

Example HAProxy configuration with one listed subnet

```
# ...
listen my-cluster-api-6443
  bind 192.168.1.100:6443
  mode tcp
  balance roundrobin
  option httpchk
  http-check connect
  http-check send meth GET uri /readyz
```

```

http-check expect status 200
  server my-cluster-master-2 192.168.1.101:6443 check inter 10s rise 2 fall 2
  server my-cluster-master-0 192.168.1.102:6443 check inter 10s rise 2 fall 2
  server my-cluster-master-1 192.168.1.103:6443 check inter 10s rise 2 fall 2

listen my-cluster-machine-config-api-22623
  bind 192.168.1.100:22623
  mode tcp
  balance roundrobin
  option httpchk
  http-check connect
  http-check send meth GET uri /healthz
  http-check expect status 200
    server my-cluster-master-2 192.168.1.101:22623 check inter 10s rise 2 fall 2
    server my-cluster-master-0 192.168.1.102:22623 check inter 10s rise 2 fall 2
    server my-cluster-master-1 192.168.1.103:22623 check inter 10s rise 2 fall 2

listen my-cluster-apps-443
  bind 192.168.1.100:443
  mode tcp
  balance roundrobin
  option httpchk
  http-check connect
  http-check send meth GET uri /healthz/ready
  http-check expect status 200
    server my-cluster-worker-0 192.168.1.111:443 check port 1936 inter 10s rise 2 fall 2
    server my-cluster-worker-1 192.168.1.112:443 check port 1936 inter 10s rise 2 fall 2
    server my-cluster-worker-2 192.168.1.113:443 check port 1936 inter 10s rise 2 fall 2

listen my-cluster-apps-80
  bind 192.168.1.100:80
  mode tcp
  balance roundrobin
  option httpchk
  http-check connect
  http-check send meth GET uri /healthz/ready
  http-check expect status 200
    server my-cluster-worker-0 192.168.1.111:80 check port 1936 inter 10s rise 2 fall 2
    server my-cluster-worker-1 192.168.1.112:80 check port 1936 inter 10s rise 2 fall 2
    server my-cluster-worker-2 192.168.1.113:80 check port 1936 inter 10s rise 2 fall 2
# ...

```

Example HAProxy configuration with multiple listed subnets

```

# ...
listen api-server-6443
  bind *:6443
  mode tcp
    server master-00 192.168.83.89:6443 check inter 1s
    server master-01 192.168.84.90:6443 check inter 1s
    server master-02 192.168.85.99:6443 check inter 1s
    server bootstrap 192.168.80.89:6443 check inter 1s

listen machine-config-server-22623
  bind *:22623
  mode tcp

```

```

server master-00 192.168.83.89:22623 check inter 1s
server master-01 192.168.84.90:22623 check inter 1s
server master-02 192.168.85.99:22623 check inter 1s
server bootstrap 192.168.80.89:22623 check inter 1s

listen ingress-router-80
  bind *:80
  mode tcp
  balance source
    server worker-00 192.168.83.100:80 check inter 1s
    server worker-01 192.168.83.101:80 check inter 1s

listen ingress-router-443
  bind *:443
  mode tcp
  balance source
    server worker-00 192.168.83.100:443 check inter 1s
    server worker-01 192.168.83.101:443 check inter 1s

listen ironic-api-6385
  bind *:6385
  mode tcp
  balance source
    server master-00 192.168.83.89:6385 check inter 1s
    server master-01 192.168.84.90:6385 check inter 1s
    server master-02 192.168.85.99:6385 check inter 1s
    server bootstrap 192.168.80.89:6385 check inter 1s

listen inspector-api-5050
  bind *:5050
  mode tcp
  balance source
    server master-00 192.168.83.89:5050 check inter 1s
    server master-01 192.168.84.90:5050 check inter 1s
    server master-02 192.168.85.99:5050 check inter 1s
    server bootstrap 192.168.80.89:5050 check inter 1s
# ...

```

2. Use the **curl** CLI command to verify that the user-managed load balancer and its resources are operational:
 - a. Verify that the cluster machine configuration API is accessible to the Kubernetes API server resource, by running the following command and observing the response:

```
$ curl https://<loadbalancer_ip_address>:6443/version --insecure
```

If the configuration is correct, you receive a JSON object in response:

```
{
  "major": "1",
  "minor": "11+",
  "gitVersion": "v1.11.0+ad103ed",
  "gitCommit": "ad103ed",
  "gitTreeState": "clean",
  "buildDate": "2019-01-09T06:44:10Z",
  "goVersion": "go1.10.3",
}
```

```
"compiler": "gc",
"platform": "linux/amd64"
}
```

- b. Verify that the cluster machine configuration API is accessible to the Machine config server resource, by running the following command and observing the output:

```
$ curl -v https://<loadbalancer_ip_address>:22623/healthz --insecure
```

If the configuration is correct, the output from the command shows the following response:

```
HTTP/1.1 200 OK
Content-Length: 0
```

- c. Verify that the controller is accessible to the Ingress Controller resource on port 80, by running the following command and observing the output:

```
$ curl -I -L -H "Host: console-openshift-console.apps.<cluster_name>.<base_domain>"
http://<load_balancer_front_end_IP_address>
```

If the configuration is correct, the output from the command shows the following response:

```
HTTP/1.1 302 Found
content-length: 0
location: https://console-openshift-console.apps.ocp4.private.opequon.net/
cache-control: no-cache
```

- d. Verify that the controller is accessible to the Ingress Controller resource on port 443, by running the following command and observing the output:

```
$ curl -I -L --insecure --resolve console-openshift-console.apps.<cluster_name>.<base_domain>:443:<Load Balancer Front End IP Address> https://console-openshift-console.apps.<cluster_name>.<base_domain>
```

If the configuration is correct, the output from the command shows the following response:


```
HTTP/1.1 200 OK
referrer-policy: strict-origin-when-cross-origin
set-cookie: csrf-
token=UIYWOyQ62LWjw2h003xtYSKlh1a0Py2hhctw0WmV2YEdhJfFyQwWcGBsja261dG
LgaYO0nxzVERhiXt6QepA7g==; Path=/; Secure; SameSite=Lax
x-content-type-options: nosniff
x-dns-prefetch-control: off
x-frame-options: DENY
x-xss-protection: 1; mode=block
date: Wed, 04 Oct 2023 16:29:38 GMT
content-type: text/html; charset=utf-8
set-cookie:
1e2670d92730b515ce3a1bb65da45062=1bf5e9573c9a2760c964ed1659cc1673; path=/;
HttpOnly; Secure; SameSite=None
cache-control: private
```

- Configure the DNS records for your cluster to target the front-end IP addresses of the user-managed load balancer. You must update records to your DNS server for the cluster API and applications over the load balancer.

Examples of modified DNS records

```
<load_balancer_ip_address> A api.<cluster_name>.<base_domain>
A record pointing to Load Balancer Front End
```

```
<load_balancer_ip_address> A apps.<cluster_name>.<base_domain>
A record pointing to Load Balancer Front End
```



IMPORTANT

DNS propagation might take some time for each DNS record to become available. Ensure that each DNS record propagates before validating each record.

- For your OpenShift Container Platform cluster to use the user-managed load balancer, you must specify the following configuration in your cluster's **install-config.yaml** file:

```
# ...
platform:
  baremetal:
    loadBalancer:
      type: UserManaged 1
      apiVIPs:
        - <api_ip> 2
      ingressVIPs:
        - <ingress_ip> 3
# ...
```

- Set **UserManaged** for the **type** parameter to specify a user-managed load balancer for your cluster. The parameter defaults to **OpenShiftManagedDefault**, which denotes the default internal load balancer. For services defined in an **openshift-kni-infra** namespace, a user-managed load balancer can deploy the **coredns** service to pods in your cluster but ignores **keepalived** and **haproxy** services.
- Required parameter when you specify a user-managed load balancer. Specify the user-managed load balancer's public IP address, so that the Kubernetes API can communicate with the user-managed load balancer.
- Required parameter when you specify a user-managed load balancer. Specify the user-managed load balancer's public IP address, so that the user-managed load balancer can manage ingress traffic for your cluster.

Verification

- Use the **curl** CLI command to verify that the user-managed load balancer and DNS record configuration are operational:
 - Verify that you can access the cluster API, by running the following command and observing the output:

```
$ curl https://api.<cluster_name>.<base_domain>:6443/version --insecure
```

If the configuration is correct, you receive a JSON object in response:

```
{
  "major": "1",
  "minor": "11+",
  "gitVersion": "v1.11.0+ad103ed",
  "gitCommit": "ad103ed",
  "gitTreeState": "clean",
  "buildDate": "2019-01-09T06:44:10Z",
  "goVersion": "go1.10.3",
  "compiler": "gc",
  "platform": "linux/amd64"
}
```

- b. Verify that you can access the cluster machine configuration, by running the following command and observing the output:

```
$ curl -v https://api.<cluster_name>.<base_domain>:22623/healthz --insecure
```

If the configuration is correct, the output from the command shows the following response:

```
HTTP/1.1 200 OK
Content-Length: 0
```

- c. Verify that you can access each cluster application on port, by running the following command and observing the output:

```
$ curl http://console-openshift-console.apps.<cluster_name>.<base_domain> -I -L --insecure
```

If the configuration is correct, the output from the command shows the following response:

```
HTTP/1.1 302 Found
content-length: 0
location: https://console-openshift-console.apps.<cluster-name>.<base domain>/
cache-control: no-cacheHTTP/1.1 200 OK
referrer-policy: strict-origin-when-cross-origin
set-cookie: csrf-
token=39HoZgztDnzjJkq/JuLJMeoKNXIfiVv2YgZc09c3TBOBU4NI6kDXaJH1LdicNhN1UsQ
Wzon4Dor9GWGfopaTEQ==; Path=/; Secure
x-content-type-options: nosniff
x-dns-prefetch-control: off
x-frame-options: DENY
x-xss-protection: 1; mode=block
date: Tue, 17 Nov 2020 08:42:10 GMT
content-type: text/html; charset=utf-8
set-cookie:
1e2670d92730b515ce3a1bb65da45062=9b714eb87e93cf34853e87a92d6894be; path=/;
HttpOnly; Secure; SameSite=None
cache-control: private
```

- d. Verify that you can access each cluster application on port 443, by running the following command and observing the output:

```
$ curl https://console-openshift-console.apps.<cluster_name>.<base_domain> -I -L --insecure
```

If the configuration is correct, the output from the command shows the following response:

```
HTTP/1.1 200 OK
referrer-policy: strict-origin-when-cross-origin
set-cookie: csrf-
token=UIYWOyQ62LWjw2h003xtYSKlh1a0Py2hhctw0WmV2YEdhJfYqWwCGBsja261dG
LgaYO0nxzVERhiXt6QepA7g==; Path=/; Secure; SameSite=Lax
x-content-type-options: nosniff
x-dns-prefetch-control: off
x-frame-options: DENY
x-xss-protection: 1; mode=block
date: Wed, 04 Oct 2023 16:29:38 GMT
content-type: text/html; charset=utf-8
set-cookie:
1e2670d92730b515ce3a1bb65da45062=1bf5e9573c9a2760c964ed1659cc1673; path=/;
HttpOnly; Secure; SameSite=None
cache-control: private
```

16.3.10. Setting the cluster node hostnames through DHCP

On Red Hat Enterprise Linux CoreOS (RHCOS) machines, hostnames are set through **NetworkManager**. By default, the machines obtain their hostnames through DHCP. If hostnames are not provided by DHCP, set statically through kernel arguments, or another method, they are obtained through a reverse DNS lookup. Reverse DNS lookup occurs after the network has been initialized on a node and can take time to resolve. Other system services can start prior to this and detect hostnames as **localhost** or similar. You can avoid this delay in assigning hostnames by using DHCP to provide the hostname for each cluster node.

Additionally, setting the hostnames through DHCP can bypass any manual DNS record name configuration errors in environments that have a DNS split-horizon implementation.

16.3.11. Configuring the `install-config.yaml` file

16.3.11.1. Configuring the `install-config.yaml` file

The **`install-config.yaml`** file requires some additional details. Most of the information teaches the installation program and the resulting cluster enough about the available hardware that it is able to fully manage it.



NOTE

The installation program no longer needs the **`clusterOSImage`** RHCOS image because the correct image is in the release payload.

1. Configure **`install-config.yaml`**. Change the appropriate variables to match the environment, including **`pullSecret`** and **`sshKey`**:


```

apiVersion: v1
baseDomain: <domain>
metadata:
  name: <cluster_name>
networking:
  machineNetwork:
    - cidr: <public_cidr>
  networkType: OVNKubernetes
compute:
- name: worker
  replicas: 2 1
controlPlane:
  name: master
  replicas: 3
platform:
  baremetal: {}
platform:
  baremetal:
    apiVIPs:
      - <api_ip>
    ingressVIPs:
      - <wildcard_ip>
    provisioningNetworkCIDR: <CIDR>
    bootstrapExternalStaticIP: <bootstrap_static_ip_address> 2
    bootstrapExternalStaticGateway: <bootstrap_static_gateway> 3
    bootstrapExternalStaticDNS: <bootstrap_static_dns> 4
  hosts:
    - name: openshift-master-0
      role: master
      bmc:
        address: ipmi://<out_of_band_ip> 5
        username: <user>
        password: <password>
        bootMACAddress: <NIC1_mac_address>
        rootDeviceHints:
          deviceName: "<installation_disk_drive_path>" 6
    - name: <openshift_master_1>
      role: master
      bmc:
        address: ipmi://<out_of_band_ip>
        username: <user>
        password: <password>
        bootMACAddress: <NIC1_mac_address>
        rootDeviceHints:
          deviceName: "<installation_disk_drive_path>"
    - name: <openshift_master_2>
      role: master
      bmc:
        address: ipmi://<out_of_band_ip>
        username: <user>
        password: <password>
        bootMACAddress: <NIC1_mac_address>
        rootDeviceHints:
          deviceName: "<installation_disk_drive_path>"
    - name: <openshift_worker_0>

```

```

role: worker
bmc:
  address: ipmi://<out_of_band_ip>
  username: <user>
  password: <password>
bootMACAddress: <NIC1_mac_address>
- name: <openshift_worker_1>
  role: worker
  bmc:
    address: ipmi://<out_of_band_ip>
    username: <user>
    password: <password>
    bootMACAddress: <NIC1_mac_address>
  rootDeviceHints:
    deviceName: "<installation_disk_drive_path>"
pullSecret: '<pull_secret>'
sshKey: '<ssh_pub_key>'

```

- 1 Scale the compute machines based on the number of compute nodes that are part of the OpenShift Container Platform cluster. Valid options for the **replicas** value are **0** and integers greater than or equal to **2**. Set the number of replicas to **0** to deploy a three-node cluster, which contains only three control plane machines. A three-node cluster is a smaller, more resource-efficient cluster that can be used for testing, development, and production. You cannot install the cluster with only one compute node.
- 2 When deploying a cluster with static IP addresses, you must set the **bootstrapExternalStaticIP** configuration setting to specify the static IP address of the bootstrap VM when there is no DHCP server on the bare-metal network.
- 3 When deploying a cluster with static IP addresses, you must set the **bootstrapExternalStaticGateway** configuration setting to specify the gateway IP address for the bootstrap VM when there is no DHCP server on the bare-metal network.
- 4 When deploying a cluster with static IP addresses, you must set the **bootstrapExternalStaticDNS** configuration setting to specify the DNS address for the bootstrap VM when there is no DHCP server on the bare-metal network.
- 5 See the BMC addressing sections for more options.
- 6 To set the path to the installation disk drive, enter the kernel name of the disk. For example, **/dev/sda**.




IMPORTANT

Because the disk discovery order is not guaranteed, the kernel name of the disk can change across booting options for machines with multiple disks. For example, `/dev/sda` becomes `/dev/sdb` and vice versa. To avoid this issue, you must use persistent disk attributes, such as the disk World Wide Name (WWN) or `/dev/disk/by-path/`. It is recommended to use the `/dev/disk/by-path/<device_path>` link to the storage location. To use the disk WWN, replace the `deviceName` parameter with the `wwnWithExtension` parameter. Depending on the parameter that you use, enter either of the following values:

- The disk name. For example, `/dev/sda`, or `/dev/disk/by-path/`.
- The disk WWN. For example, `"0x64cd98f04fde100024684cf3034da5c2"`. Ensure that you enter the disk WWN value within quotes so that it is used as a string value and not a hexadecimal value.

Failure to meet these requirements for the `rootDeviceHints` parameter might result in the following error:

```
ironic-inspector inspection failed: No disks satisfied root device hints
```



NOTE

Before OpenShift Container Platform 4.12, the cluster installation program only accepted an IPv4 address or an IPv6 address for the `apiVIP` and `ingressVIP` configuration settings. In OpenShift Container Platform 4.12 and later, these configuration settings are deprecated. Instead, use a list format in the `apiVIPs` and `ingressVIPs` configuration settings to specify IPv4 addresses, IPv6 addresses, or both IP address formats.

2. Create a directory to store the cluster configuration:

```
$ mkdir ~/clusterconfigs
```

3. Copy the `install-config.yaml` file to the new directory:

```
$ cp install-config.yaml ~/clusterconfigs
```

4. Ensure all bare metal nodes are powered off prior to installing the OpenShift Container Platform cluster:

```
$ ipmitool -I lanplus -U <user> -P <password> -H <management-server-ip> power off
```

5. Remove old bootstrap resources if any are left over from a previous deployment attempt:

```
for i in $(sudo virsh list | tail -n +3 | grep bootstrap | awk {'print $2'});
do
  sudo virsh destroy $i;
  sudo virsh undefine $i;
```

```

sudo virsh vol-delete $i --pool $i;
sudo virsh vol-delete $i.ign --pool $i;
sudo virsh pool-destroy $i;
sudo virsh pool-undefine $i;
done

```


16.3.11.2. Additional install-config parameters

See the following tables for the required parameters, the **hosts** parameter, and the **bmc** parameter for the **install-config.yaml** file.

Table 16.5. Required parameters

Parameters	Default	Description
baseDomain		The domain name for the cluster. For example, example.com .
bootMode	UEFI	The boot mode for a node. Options are legacy , UEFI , and UEFISecureBoot . If bootMode is not set, Ironic sets it while inspecting the node.
bootstrapExternalStaticDNS		The static network DNS of the bootstrap node. You must set this value when deploying a cluster with static IP addresses when there is no Dynamic Host Configuration Protocol (DHCP) server on the bare-metal network. If you do not set this value, the installation program will use the value from bootstrapExternalStaticGateway , which causes problems when the IP address values of the gateway and DNS are different.
bootstrapExternalStaticIP		The static IP address for the bootstrap VM. You must set this value when deploying a cluster with static IP addresses when there is no DHCP server on the bare-metal network.
bootstrapExternalStaticGateway		The static IP address of the gateway for the bootstrap VM. You must set this value when deploying a cluster with static IP addresses when there is no DHCP server on the bare-metal network.
sshKey		The sshKey configuration setting contains the key in the <code>~/.ssh/id_rsa.pub</code> file required to access the control plane nodes and compute nodes. Typically, this key is from the provisioner node.
pullSecret		The pullSecret configuration setting contains a copy of the pull secret downloaded from the Install OpenShift on Bare Metal page when preparing the provisioner node.
metadata: name:		The name to be given to the OpenShift Container Platform cluster. For example, openshift .

Parameters	Default	Description
networking: machineNetwork: - cidr:		The public CIDR (Classless Inter-Domain Routing) of the external network. For example, 10.0.0.0/24 .
compute: - name: worker		The OpenShift Container Platform cluster requires a name be provided for compute nodes even if there are zero nodes.
compute: replicas: 2		Replicas sets the number of compute nodes in the OpenShift Container Platform cluster.
controlPlane: name: master		The OpenShift Container Platform cluster requires a name for control plane nodes.
controlPlane: replicas: 3		Replicas sets the number of control plane nodes included as part of the OpenShift Container Platform cluster.
provisioningNetwork Interface		The name of the network interface on nodes connected to the provisioning network. For OpenShift Container Platform 4.9 and later releases, use the bootMACAddress configuration setting to enable Ironic to identify the IP address of the NIC instead of using the provisioningNetworkInterface configuration setting to identify the name of the NIC.
defaultMachinePlatform		The default configuration used for machine pools without a platform configuration.

Parameters	Default	Description
apiVIPs		<p>(Optional) The virtual IP address for Kubernetes API communication.</p> <p>This setting must either be provided in the install-config.yaml file as a reserved IP from the MachineNetwork or preconfigured in the DNS so that the default name resolves correctly. Use the virtual IP address and not the FQDN when adding a value to the apiVIPs configuration setting in the install-config.yaml file. The primary IP address must be from the IPv4 network when using dual stack networking. If not set, the installation program uses api.<cluster_name>.<base_domain> to derive the IP address from the DNS.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p>NOTE</p> <p>Before OpenShift Container Platform 4.12, the cluster installation program only accepted an IPv4 address or an IPv6 address for the apiVIP configuration setting. From OpenShift Container Platform 4.12 or later, the apiVIP configuration setting is deprecated. Instead, use a list format for the apiVIPs configuration setting to specify an IPv4 address, an IPv6 address or both IP address formats.</p> </div> </div>
disableCertificateVerification	False	redfish and redfish-virtualmedia need this parameter to manage BMC addresses. The value should be True when using a self-signed certificate for BMC addresses.

Parameters	Default	Description
ingressVIPs		<p>(Optional) The virtual IP address for ingress traffic.</p> <p>This setting must either be provided in the install-config.yaml file as a reserved IP from the MachineNetwork or preconfigured in the DNS so that the default name resolves correctly. Use the virtual IP address and not the FQDN when adding a value to the ingressVIPs configuration setting in the install-config.yaml file. The primary IP address must be from the IPv4 network when using dual stack networking. If not set, the installation program uses test.apps.<cluster_name>.<base_domain> to derive the IP address from the DNS.</p> <div style="display: flex; align-items: flex-start;"> <div style="width: 40px; height: 100px; background: repeating-linear-gradient(45deg, transparent, transparent 2px, #ccc 2px, #ccc 4px); border: 1px solid #ccc; margin-right: 10px;"></div> <div> <p>NOTE</p> <p>Before OpenShift Container Platform 4.12, the cluster installation program only accepted an IPv4 address or an IPv6 address for the ingressVIP configuration setting. In OpenShift Container Platform 4.12 and later, the ingressVIP configuration setting is deprecated. Instead, use a list format for the ingressVIPs configuration setting to specify an IPv4 addresses, an IPv6 addresses or both IP address formats.</p> </div> </div>

Table 16.6. Optional Parameters


Parameters	Default	Description
provisioningDHCPRange	172.22.0.10,172.22.0.100	Defines the IP range for nodes on the provisioning network.
provisioningNetworkCIDR	172.22.0.0/24	The CIDR for the network to use for provisioning. This option is required when not using the default address range on the provisioning network.
clusterProvisioningIP	The third IP address of the provisioningNetworkCIDR .	The IP address within the cluster where the provisioning services run. Defaults to the third IP address of the provisioning subnet. For example, 172.22.0.3 .
bootstrapProvisioningIP	The second IP address of the provisioningNetworkCIDR .	The IP address on the bootstrap VM where the provisioning services run while the installer is deploying the control plane (master) nodes. Defaults to the second IP address of the provisioning subnet. For example, 172.22.0.2 or 2620:52:0:1307::2 .
externalBridge	baremetal	The name of the bare-metal bridge of the hypervisor attached to the bare-metal network.

Parameters	Default	Description
provisioningBridge	provisioning	The name of the provisioning bridge on the provisioner host attached to the provisioning network.
architecture		Defines the host architecture for your cluster. Valid values are amd64 or arm64 .
defaultMachinePlatform		The default configuration used for machine pools without a platform configuration.
bootstrapOSImage		A URL to override the default operating system image for the bootstrap node. The URL must contain a SHA-256 hash of the image. For example: <a href="https://mirror.openshift.com/rhcos-<version>-qemu.qcow2.gz?sha256=<uncompressed_sha256>">https://mirror.openshift.com/rhcos-<version>-qemu.qcow2.gz?sha256=<uncompressed_sha256> .
provisioningNetwork		<p>The provisioningNetwork configuration setting determines whether the cluster uses the provisioning network. If it does, the configuration setting also determines if the cluster manages the network.</p> <p>Disabled: Set this parameter to Disabled to disable the requirement for a provisioning network. When set to Disabled, you must only use virtual media based provisioning, or bring up the cluster using the assisted installer. If Disabled and using power management, BMCs must be accessible from the bare-metal network. If Disabled, you must provide two IP addresses on the bare-metal network that are used for the provisioning services.</p> <p>Managed: Set this parameter to Managed, which is the default, to fully manage the provisioning network, including DHCP, TFTP, and so on.</p> <p>Unmanaged: Set this parameter to Unmanaged to enable the provisioning network but take care of manual configuration of DHCP. Virtual media provisioning is recommended but PXE is still available if required.</p>
httpProxy		Set this parameter to the appropriate HTTP proxy used within your environment.
httpsProxy		Set this parameter to the appropriate HTTPS proxy used within your environment.
noProxy		Set this parameter to the appropriate list of exclusions for proxy usage within your environment.

Hosts

The **hosts** parameter is a list of separate bare metal assets used to build the cluster.

Table 16.7. Hosts

Name	Default	Description
name		The name of the BareMetalHost resource to associate with the details. For example, openshift-master-0 .
role		The role of the bare metal node. Either master (control plane node) or worker (compute node).
bmc		Connection details for the baseboard management controller. See the BMC addressing section for additional details.
bootMACAddress		<p>The MAC address of the NIC that the host uses for the provisioning network. Ironic retrieves the IP address using the bootMACAddress configuration setting. Then, it binds to the host.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p>NOTE</p> <p>You must provide a valid MAC address from the host if you disabled the provisioning network.</p> </div> </div>
networkConfig		Set this optional parameter to configure the network interface of a host. See "(Optional) Configuring host network interfaces" for additional details.

16.3.11.3. BMC addressing

Most vendors support Baseboard Management Controller (BMC) addressing with the Intelligent Platform Management Interface (IPMI). IPMI does not encrypt communications. It is suitable for use within a data center over a secured or dedicated management network. Check with your vendor to see if they support Redfish network boot. Redfish delivers simple and secure management for converged, hybrid IT and the Software Defined Data Center (SDDC). Redfish is human readable and machine capable, and leverages common internet and web services standards to expose information directly to the modern tool chain. If your hardware does not support Redfish network boot, use IPMI.

You can modify the BMC address during installation while the node is in the **Registering** state. If you need to modify the BMC address after the node leaves the **Registering** state, you must disconnect the node from Ironic, edit the **BareMetalHost** resource, and reconnect the node to Ironic. See the *Editing a BareMetalHost resource* section for details.

IPMI

Hosts using IPMI use the **ipmi://<out-of-band-ip>:<port>** address format, which defaults to port **623** if not specified. The following example demonstrates an IPMI configuration within the **install-config.yaml** file.

```
platform:
  baremetal:
```

```
hosts:
- name: openshift-master-0
  role: master
  bmc:
    address: ipmi://<out-of-band-ip>
    username: <user>
    password: <password>
```



IMPORTANT

The **provisioning** network is required when PXE booting using IPMI for BMC addressing. It is not possible to PXE boot hosts without a **provisioning** network. If you deploy without a **provisioning** network, you must use a virtual media BMC addressing option such as **redfish-virtualmedia** or **idrac-virtualmedia**. See "Redfish virtual media for HPE iLO" in the "BMC addressing for HPE iLO" section or "Redfish virtual media for Dell iDRAC" in the "BMC addressing for Dell iDRAC" section for additional details.

Redfish network boot

To enable Redfish, use **redfish://** or **redfish+http://** to disable TLS. The installer requires both the hostname or the IP address and the path to the system ID. The following example demonstrates a Redfish configuration within the **install-config.yaml** file.

```
platform:
  baremetal:
    hosts:
      - name: openshift-master-0
        role: master
        bmc:
          address: redfish://<out-of-band-ip>/redfish/v1/Systems/1
          username: <user>
          password: <password>
```

While it is recommended to have a certificate of authority for the out-of-band management addresses, you must include **disableCertificateVerification: True** in the **bmc** configuration if using self-signed certificates. The following example demonstrates a Redfish configuration using the **disableCertificateVerification: True** configuration parameter within the **install-config.yaml** file.

```
platform:
  baremetal:
    hosts:
      - name: openshift-master-0
        role: master
        bmc:
          address: redfish://<out-of-band-ip>/redfish/v1/Systems/1
          username: <user>
          password: <password>
          disableCertificateVerification: True
```

Additional resources

- [Editing a BareMetalHost resource](#)

16.3.11.4. Verifying support for Redfish APIs

When installing using the Redfish API, the installation program calls several Redfish endpoints on the baseboard management controller (BMC) when using installer-provisioned infrastructure on bare metal. If you use Redfish, ensure that your BMC supports all of the Redfish APIs before installation.

Procedure

1. Set the IP address or hostname of the BMC by running the following command:

```
$ export SERVER=<ip_address> 1
```

- 1 Replace **<ip_address>** with the IP address or hostname of the BMC.

2. Set the ID of the system by running the following command:

```
$ export SystemID=<system_id> 1
```

- 1 Replace **<system_id>** with the system ID. For example, **System.Embedded.1** or **1**. See the following vendor-specific BMC sections for details.

List of Redfish APIs:

1. Check **power on** support by running the following command:

```
$ curl -u $USER:$PASS -X POST -H'Content-Type: application/json' -H'Accept: application/json' -d '{"ResetType": "On"}' https://$SERVER/redfish/v1/Systems/$SystemID/Actions/ComputerSystem.Reset
```

2. Check **power off** support by running the following command:

```
$ curl -u $USER:$PASS -X POST -H'Content-Type: application/json' -H'Accept: application/json' -d '{"ResetType": "ForceOff"}' https://$SERVER/redfish/v1/Systems/$SystemID/Actions/ComputerSystem.Reset
```

3. Check the temporary boot implementation that uses **pxe** by running the following command:

```
$ curl -u $USER:$PASS -X PATCH -H "Content-Type: application/json" https://$Server/redfish/v1/Systems/$SystemID/ -d '{"Boot": {"BootSourceOverrideTarget": "pxe", "BootSourceOverrideEnabled": "Once"}}'
```

4. Check the status of setting the BIOS boot mode that uses **Legacy** or **UEFI** by running the following command:

```
$ curl -u $USER:$PASS -X PATCH -H "Content-Type: application/json" https://$Server/redfish/v1/Systems/$SystemID/ -d '{"Boot": {"BootSourceOverrideMode": "UEFI"}}'
```

List of Redfish virtual media APIs:

1. Check the ability to set the temporary boot device that uses **cd** or **dvd** by running the following command:

```
$ curl -u $USER:$PASS -X PATCH -H "Content-Type: application/json"
https://$Server/redfish/v1/Systems/$SystemID/ -d '{"Boot": {"BootSourceOverrideTarget":
"cd", "BootSourceOverrideEnabled": "Once"}}'
```

2. Check the ability to mount virtual media by running the following command:

```
$ curl -u $USER:$PASS -X PATCH -H "Content-Type: application/json" -H "If-Match: *"
https://$Server/redfish/v1/Managers/$ManagerID/VirtualMedia/$VmediaId -d '{"Image":
"https://example.com/test.iso", "TransferProtocolType": "HTTPS", "UserName": "",
"Password":""}'
```



NOTE

The **PowerOn** and **PowerOff** commands for Redfish APIs are the same for the Redfish virtual media APIs.



IMPORTANT

HTTPS and **HTTP** are the only supported parameter types for **TransferProtocolTypes**.

16.3.11.5. BMC addressing for Dell iDRAC

The **address** field for each **bmc** entry is a URL for connecting to the OpenShift Container Platform cluster nodes, including the type of controller in the URL scheme and its location on the network.

```
platform:
  baremetal:
    hosts:
      - name: <hostname>
        role: <master | worker>
      bmc:
        address: <address> 1
        username: <user>
        password: <password>
```

- 1** The **address** configuration setting specifies the protocol.

For Dell hardware, Red Hat supports integrated Dell Remote Access Controller (iDRAC) virtual media, Redfish network boot, and IPMI.

BMC address formats for Dell iDRAC

Protocol	Address Format
iDRAC virtual media	idrac-virtualmedia://<out-of-band-ip>/redfish/v1/Systems/System.Embedded.1
Redfish network boot	redfish://<out-of-band-ip>/redfish/v1/Systems/System.Embedded.1
IPMI	ipmi://<out-of-band-ip>



IMPORTANT

Use **idrac-virtualmedia** as the protocol for Redfish virtual media. **redfish-virtualmedia** will not work on Dell hardware. Dell's **idrac-virtualmedia** uses the Redfish standard with Dell's OEM extensions.

See the following sections for additional details.

Redfish virtual media for Dell iDRAC

For Redfish virtual media on Dell servers, use **idrac-virtualmedia://** in the **address** setting. Using **redfish-virtualmedia://** will not work.



NOTE

Use **idrac-virtualmedia://** as the protocol for Redfish virtual media. Using **redfish-virtualmedia://** will not work on Dell hardware, because the **idrac-virtualmedia://** protocol corresponds to the **idrac** hardware type and the Redfish protocol in Ironic. Dell's **idrac-virtualmedia://** protocol uses the Redfish standard with Dell's OEM extensions. Ironic also supports the **idrac** type with the WSMAN protocol. Therefore, you must specify **idrac-virtualmedia://** to avoid unexpected behavior when electing to use Redfish with virtual media on Dell hardware.

The following example demonstrates using iDRAC virtual media within the **install-config.yaml** file.

```

platform:
  baremetal:
    hosts:
      - name: openshift-master-0
        role: master
    bmc:
      address: idrac-virtualmedia://<out-of-band-ip>/redfish/v1/Systems/System.Embedded.1
      username: <user>
      password: <password>

```

While it is recommended to have a certificate of authority for the out-of-band management addresses, you must include **disableCertificateVerification: True** in the **bmc** configuration if using self-signed certificates.



NOTE

Ensure the OpenShift Container Platform cluster nodes have **AutoAttach** enabled through the iDRAC console. The menu path is: **Configuration** → **Virtual Media** → **Attach Mode** → **AutoAttach**.

The following example demonstrates a Redfish configuration using the **disableCertificateVerification: True** configuration parameter within the **install-config.yaml** file.

```

platform:
  baremetal:
    hosts:
      - name: openshift-master-0
        role: master
    bmc:
      address: idrac-virtualmedia://<out-of-band-ip>/redfish/v1/Systems/System.Embedded.1

```

```
username: <user>
password: <password>
disableCertificateVerification: True
```

Redfish network boot for iDRAC

To enable Redfish, use **redfish://** or **redfish+http://** to disable transport layer security (TLS). The installer requires both the hostname or the IP address and the path to the system ID. The following example demonstrates a Redfish configuration within the **install-config.yaml** file.

```
platform:
  baremetal:
    hosts:
      - name: openshift-master-0
        role: master
    bmc:
      address: redfish://<out-of-band-ip>/redfish/v1/Systems/System.Embedded.1
      username: <user>
      password: <password>
```

While it is recommended to have a certificate of authority for the out-of-band management addresses, you must include **disableCertificateVerification: True** in the **bmc** configuration if using self-signed certificates. The following example demonstrates a Redfish configuration using the **disableCertificateVerification: True** configuration parameter within the **install-config.yaml** file.

```
platform:
  baremetal:
    hosts:
      - name: openshift-master-0
        role: master
    bmc:
      address: redfish://<out-of-band-ip>/redfish/v1/Systems/System.Embedded.1
      username: <user>
      password: <password>
      disableCertificateVerification: True
```

NOTE

There is a known issue on Dell iDRAC 9 with firmware version **04.40.00.00** and all releases up to including the **5.xx** series for installer-provisioned installations on bare metal deployments. The virtual console plugin defaults to eHTML5, an enhanced version of HTML5, which causes problems with the **InsertVirtualMedia** workflow. Set the plugin to use HTML5 to avoid this issue. The menu path is **Configuration → Virtual console → Plug-in Type → HTML5**.

Ensure the OpenShift Container Platform cluster nodes have **AutoAttach** enabled through the iDRAC console. The menu path is: **Configuration → Virtual Media → Attach Mode → AutoAttach**.

16.3.11.6. BMC addressing for HPE iLO

The **address** field for each **bmc** entry is a URL for connecting to the OpenShift Container Platform cluster nodes, including the type of controller in the URL scheme and its location on the network.

```
platform:
```

```

baremetal:
  hosts:
    - name: <hostname>
      role: <master | worker>
      bmc:
        address: <address> ❶
        username: <user>
        password: <password>

```

❶ The **address** configuration setting specifies the protocol.

For HPE integrated Lights Out (iLO), Red Hat supports Redfish virtual media, Redfish network boot, and IPMI.

Table 16.8. BMC address formats for HPE iLO

Protocol	Address Format
Redfish virtual media	redfish-virtualmedia://<out-of-band-ip>/redfish/v1/Systems/1
Redfish network boot	redfish://<out-of-band-ip>/redfish/v1/Systems/1
IPMI	ipmi://<out-of-band-ip>

See the following sections for additional details.

Redfish virtual media for HPE iLO

To enable Redfish virtual media for HPE servers, use **redfish-virtualmedia://** in the **address** setting. The following example demonstrates using Redfish virtual media within the **install-config.yaml** file.

```

platform:
  baremetal:
    hosts:
      - name: openshift-master-0
        role: master
        bmc:
          address: redfish-virtualmedia://<out-of-band-ip>/redfish/v1/Systems/1
          username: <user>
          password: <password>

```

While it is recommended to have a certificate of authority for the out-of-band management addresses, you must include **disableCertificateVerification: True** in the **bmc** configuration if using self-signed certificates. The following example demonstrates a Redfish configuration using the **disableCertificateVerification: True** configuration parameter within the **install-config.yaml** file.

```

platform:
  baremetal:
    hosts:
      - name: openshift-master-0
        role: master
        bmc:
          address: redfish-virtualmedia://<out-of-band-ip>/redfish/v1/Systems/1

```

```
username: <user>
password: <password>
disableCertificateVerification: True
```



NOTE

Redfish virtual media is not supported on 9th generation systems running iLO4, because Ironic does not support iLO4 with virtual media.

Redfish network boot for HPE iLO

To enable Redfish, use **redfish://** or **redfish+http://** to disable TLS. The installer requires both the hostname or the IP address and the path to the system ID. The following example demonstrates a Redfish configuration within the **install-config.yaml** file.

```
platform:
  baremetal:
    hosts:
      - name: openshift-master-0
        role: master
    bmc:
      address: redfish://<out-of-band-ip>/redfish/v1/Systems/1
      username: <user>
      password: <password>
```

While it is recommended to have a certificate of authority for the out-of-band management addresses, you must include **disableCertificateVerification: True** in the **bmc** configuration if using self-signed certificates. The following example demonstrates a Redfish configuration using the **disableCertificateVerification: True** configuration parameter within the **install-config.yaml** file.

```
platform:
  baremetal:
    hosts:
      - name: openshift-master-0
        role: master
    bmc:
      address: redfish://<out-of-band-ip>/redfish/v1/Systems/1
      username: <user>
      password: <password>
      disableCertificateVerification: True
```

16.3.11.7. BMC addressing for Fujitsu iRMC

The **address** field for each **bmc** entry is a URL for connecting to the OpenShift Container Platform cluster nodes, including the type of controller in the URL scheme and its location on the network.

```
platform:
  baremetal:
    hosts:
      - name: <hostname>
        role: <master | worker>
    bmc:
```



```
address: <address> 1
username: <user>
password: <password>
```

- 1** The **address** configuration setting specifies the protocol.

For Fujitsu hardware, Red Hat supports integrated Remote Management Controller (iRMC) and IPMI.

Table 16.9. BMC address formats for Fujitsu iRMC

Protocol	Address Format
iRMC	irmc://<out-of-band-ip>
IPMI	ipmi://<out-of-band-ip>

iRMC

Fujitsu nodes can use **irmc://<out-of-band-ip>** and defaults to port **443**. The following example demonstrates an iRMC configuration within the **install-config.yaml** file.

```
platform:
  baremetal:
    hosts:
      - name: openshift-master-0
        role: master
    bmc:
      address: irmc://<out-of-band-ip>
      username: <user>
      password: <password>
```



NOTE

Currently Fujitsu supports iRMC S5 firmware version 3.05P and above for installer-provisioned installation on bare metal.

16.3.11.8. Root device hints

The **rootDeviceHints** parameter enables the installer to provision the Red Hat Enterprise Linux CoreOS (RHCOS) image to a particular device. The installer examines the devices in the order it discovers them, and compares the discovered values with the hint values. The installer uses the first discovered device that matches the hint value. The configuration can combine multiple hints, but a device must match all hints for the installer to select it.

Table 16.10. Subfields

Subfield	Description
----------	-------------

Subfield	Description
deviceName	A string containing a Linux device name such as /dev/vda or /dev/disk/by-path/ . It is recommended to use the /dev/disk/by-path/<device_path> link to the storage location. The hint must match the actual value exactly.
hctl	A string containing a SCSI bus address like 0:0:0:0 . The hint must match the actual value exactly.
model	A string containing a vendor-specific device identifier. The hint can be a substring of the actual value.
vendor	A string containing the name of the vendor or manufacturer of the device. The hint can be a substring of the actual value.
serialNumber	A string containing the device serial number. The hint must match the actual value exactly.
minSizeGigabytes	An integer representing the minimum size of the device in gigabytes.
wwn	A string containing the unique storage identifier. The hint must match the actual value exactly.
wwnWithExtension	A string containing the unique storage identifier with the vendor extension appended. The hint must match the actual value exactly.
wwnVendorExtension	A string containing the unique vendor storage identifier. The hint must match the actual value exactly.
rotational	A boolean indicating whether the device should be a rotating disk (true) or not (false).

Example usage

```
- name: master-0
  role: master
  bmc:
    address: ipmi://10.10.0.3:6203
    username: admin
    password: redhat
```

```
bootMACAddress: de:ad:be:ef:00:40
rootDeviceHints:
  deviceName: "/dev/sda"
```

16.3.11.9. Optional: Setting proxy settings

To deploy an OpenShift Container Platform cluster using a proxy, make the following changes to the **install-config.yaml** file.

```
apiVersion: v1
baseDomain: <domain>
proxy:
  httpProxy: http://USERNAME:PASSWORD@proxy.example.com:PORT
  httpsProxy: https://USERNAME:PASSWORD@proxy.example.com:PORT
  noProxy: <WILDCARD_OF_DOMAIN>,<PROVISIONING_NETWORK/CIDR>,
  <BMC_ADDRESS_RANGE/CIDR>
```

The following is an example of **noProxy** with values.

```
noProxy: .example.com,172.22.0.0/24,10.10.0.0/24
```

With a proxy enabled, set the appropriate values of the proxy in the corresponding key/value pair.

Key considerations:

- If the proxy does not have an HTTPS proxy, change the value of **httpsProxy** from **https://** to **http://**.
- If using a provisioning network, include it in the **noProxy** setting, otherwise the installer will fail.
- Set all of the proxy settings as environment variables within the provisioner node. For example, **HTTP_PROXY**, **HTTPS_PROXY**, and **NO_PROXY**.



NOTE

When provisioning with IPv6, you cannot define a CIDR address block in the **noProxy** settings. You must define each address separately.

16.3.11.10. Optional: Deploying with no provisioning network

To deploy an OpenShift Container Platform cluster without a **provisioning** network, make the following changes to the **install-config.yaml** file.

```
platform:
  baremetal:
    apiVIPs:
      - <api_VIP>
    ingressVIPs:
      - <ingress_VIP>
    provisioningNetwork: "Disabled" 1
```

- 1** Add the **provisioningNetwork** configuration setting, if needed, and set it to **Disabled**.

**IMPORTANT**

The **provisioning** network is required for PXE booting. If you deploy without a **provisioning** network, you must use a virtual media BMC addressing option such as **redfish-virtualmedia** or **idrac-virtualmedia**. See "Redfish virtual media for HPE iLO" in the "BMC addressing for HPE iLO" section or "Redfish virtual media for Dell iDRAC" in the "BMC addressing for Dell iDRAC" section for additional details.

16.3.11.11. Optional: Deploying with dual-stack networking

For dual-stack networking in OpenShift Container Platform clusters, you can configure IPv4 and IPv6 address endpoints for cluster nodes. To configure IPv4 and IPv6 address endpoints for cluster nodes, edit the **machineNetwork**, **clusterNetwork**, and **serviceNetwork** configuration settings in the **install-config.yaml** file. Each setting must have two CIDR entries each. For a cluster with the IPv4 family as the primary address family, specify the IPv4 setting first. For a cluster with the IPv6 family as the primary address family, specify the IPv6 setting first.

```
machineNetwork:
- cidr: {{ extcidrnet }}
- cidr: {{ extcidrnet6 }}
clusterNetwork:
- cidr: 10.128.0.0/14
  hostPrefix: 23
- cidr: fd02::/48
  hostPrefix: 64
serviceNetwork:
- 172.30.0.0/16
- fd03::/112
```

**IMPORTANT**

On a bare-metal platform, if you specified an NMState configuration in the **networkConfig** section of your **install-config.yaml** file, add **interfaces.wait-ip: ipv4+ipv6** to the NMState YAML file to resolve an issue that prevents your cluster from deploying on a dual-stack network.

Example NMState YAML configuration file that includes the wait-ip parameter

```
networkConfig:
  nmstate:
    interfaces:
      - name: <interface_name>
        # ...
        wait-ip: ipv4+ipv6
        # ...
```

To provide an interface to the cluster for applications that use IPv4 and IPv6 addresses, configure IPv4 and IPv6 virtual IP (VIP) address endpoints for the Ingress VIP and API VIP services. To configure IPv4 and IPv6 address endpoints, edit the **apiVIPs** and **ingressVIPs** configuration settings in the **install-config.yaml** file. The **apiVIPs** and **ingressVIPs** configuration settings use a list format. The order of the list indicates the primary and secondary VIP address for each service.

```
platform:
```

```

baremetal:
  apiVIPs:
    - <api_ipv4>
    - <api_ipv6>
  ingressVIPs:
    - <wildcard_ipv4>
    - <wildcard_ipv6>

```

**NOTE**

For a cluster with dual-stack networking configuration, you must assign both IPv4 and IPv6 addresses to the same interface.

16.3.11.12. Optional: Configuring host network interfaces

Before installation, you can set the **networkConfig** configuration setting in the **install-config.yaml** file to configure host network interfaces using NMState.

The most common use case for this functionality is to specify a static IP address on the bare-metal network, but you can also configure other networks such as a storage network. This functionality supports other NMState features such as VLAN, VXLAN, bridges, bonds, routes, MTU, and DNS resolver settings.

Prerequisites

- Configure a **PTR** DNS record with a valid hostname for each node with a static IP address.
- Install the NMState CLI (**nmstate**).

Procedure

1. Optional: Consider testing the NMState syntax with **nmstatectl gc** before including it in the **install-config.yaml** file, because the installer will not check the NMState YAML syntax.

**NOTE**

Errors in the YAML syntax might result in a failure to apply the network configuration. Additionally, maintaining the validated YAML syntax is useful when applying changes using Kubernetes NMState after deployment or when expanding the cluster.

- a. Create an NMState YAML file:

```

interfaces:
  - name: <nic1_name> 1
    type: ethernet
    state: up
    ipv4:
      address:
        - ip: <ip_address> 2
          prefix-length: 24
        enabled: true
    dns-resolver:
      config:

```

```

server:
- <dns_ip_address> 3
routes:
config:
- destination: 0.0.0.0/0
  next-hop-address: <next_hop_ip_address> 4
  next-hop-interface: <next_hop_nic1_name> 5

```

1 2 3 4 5 Replace **<nic1_name>**, **<ip_address>**, **<dns_ip_address>**, **<next_hop_ip_address>** and **<next_hop_nic1_name>** with appropriate values.

- b. Test the configuration file by running the following command:

```
$ nmstatectl gc <nmstate_yaml_file>
```

Replace **<nmstate_yaml_file>** with the configuration file name.

2. Use the **networkConfig** configuration setting by adding the NMState configuration to hosts within the **install-config.yaml** file:

```

hosts:
- name: openshift-master-0
  role: master
  bmc:
    address: redfish+http://<out_of_band_ip>/redfish/v1/Systems/
    username: <user>
    password: <password>
    disableCertificateVerification: null
  bootMACAddress: <NIC1_mac_address>
  bootMode: UEFI
  rootDeviceHints:
    deviceName: "/dev/sda"
  networkConfig: 1
  interfaces:
  - name: <nic1_name> 2
    type: ethernet
    state: up
    ipv4:
      address:
      - ip: <ip_address> 3
        prefix-length: 24
      enabled: true
  dns-resolver:
    config:
      server:
      - <dns_ip_address> 4
  routes:
    config:
      - destination: 0.0.0.0/0
        next-hop-address: <next_hop_ip_address> 5
        next-hop-interface: <next_hop_nic1_name> 6

```

- 1 Add the NMState YAML syntax to configure the host interfaces.
- 2 3 4 5 6 Replace `<nic1_name>`, `<ip_address>`, `<dns_ip_address>`, `<next_hop_ip_address>` and `<next_hop_nic1_name>` with appropriate values.



IMPORTANT

After deploying the cluster, you cannot modify the **networkConfig** configuration setting of **install-config.yaml** file to make changes to the host network interface. Use the Kubernetes NMState Operator to make changes to the host network interface after deployment.

16.3.11.13. Configuring host network interfaces for subnets

For edge computing scenarios, it can be beneficial to locate compute nodes closer to the edge. To locate remote nodes in subnets, you might use different network segments or subnets for the remote nodes than you used for the control plane subnet and local compute nodes. You can reduce latency for the edge and allow for enhanced scalability by setting up subnets for edge computing scenarios.



IMPORTANT

When using the default load balancer, **OpenShiftManagedDefault** and adding remote nodes to your OpenShift Container Platform cluster, all control plane nodes must run in the same subnet. When using more than one subnet, you can also configure the Ingress VIP to run on the control plane nodes by using a manifest. See "Configuring network components to run on the control plane" for details.

If you have established different network segments or subnets for remote nodes as described in the section on "Establishing communication between subnets", you must specify the subnets in the **machineNetwork** configuration setting if the workers are using static IP addresses, bonds or other advanced networking. When setting the node IP address in the **networkConfig** parameter for each remote node, you must also specify the gateway and the DNS server for the subnet containing the control plane nodes when using static IP addresses. This ensures the remote nodes can reach the subnet containing the control plane nodes and that they can receive network traffic from the control plane.



NOTE

Deploying a cluster with multiple subnets requires using virtual media, such as **redfish-virtualmedia** and **idrac-virtualmedia**, because remote nodes cannot access the local provisioning network.

Procedure

1. Add the subnets to the **machineNetwork** in the **install-config.yaml** file when using static IP addresses:

```
networking:
  machineNetwork:
    - cidr: 10.0.0.0/24
    - cidr: 192.168.0.0/24
  networkType: OVNKubernetes
```

2. Add the gateway and DNS configuration to the **networkConfig** parameter of each edge compute node using NMState syntax when using a static IP address or advanced networking such as bonds:

```
networkConfig:
  interfaces:
  - name: <interface_name> ❶
    type: ethernet
    state: up
    ipv4:
      enabled: true
      dhcp: false
      address:
      - ip: <node_ip> ❷
        prefix-length: 24
      gateway: <gateway_ip> ❸
    dns-resolver:
      config:
        server:
        - <dns_ip> ❹
```

- ❶ Replace **<interface_name>** with the interface name.
- ❷ Replace **<node_ip>** with the IP address of the node.
- ❸ Replace **<gateway_ip>** with the IP address of the gateway.
- ❹ Replace **<dns_ip>** with the IP address of the DNS server.

16.3.11.14. Optional: Configuring address generation modes for SLAAC in dual-stack networks

For dual-stack clusters that use Stateless Address AutoConfiguration (SLAAC), you must specify a global value for the **ipv6.addr-gen-mode** network setting. You can set this value using NMState to configure the RAM disk and the cluster configuration files. If you do not configure a consistent **ipv6.addr-gen-mode** in these locations, IPv6 address mismatches can occur between CSR resources and **BareMetalHost** resources in the cluster.

Prerequisites

- Install the NMState CLI (**nmstate**).

Procedure

1. Optional: Consider testing the NMState YAML syntax with the **nmstatectl gc** command before including it in the **install-config.yaml** file because the installation program will not check the NMState YAML syntax.
 - a. Create an NMState YAML file:

```
interfaces:
- name: eth0
  ipv6:
    addr-gen-mode: <address_mode> ❶
```


-
- 1 Replace **<address_mode>** with the type of address generation mode required for IPv6 addresses in the cluster. Valid values are **eui64**, **stable-privacy**, or **random**.

b. Test the configuration file by running the following command:

```
$ nmstatectl gc <nmstate_yaml_file> 1
```

- 1 Replace **<nmstate_yaml_file>** with the name of the test configuration file.

2. Add the NMState configuration to the **hosts.networkConfig** section within the `install-config.yaml` file:

```
hosts:
  - name: openshift-master-0
    role: master
    bmc:
      address: redfish+http://<out_of_band_ip>/redfish/v1/Systems/
      username: <user>
      password: <password>
      disableCertificateVerification: null
    bootMACAddress: <NIC1_mac_address>
    bootMode: UEFI
    rootDeviceHints:
      deviceName: "/dev/sda"
    networkConfig:
      interfaces:
        - name: eth0
          ipv6:
            addr-gen-mode: <address_mode> 1
  ...
```

- 1 Replace **<address_mode>** with the type of address generation mode required for IPv6 addresses in the cluster. Valid values are **eui64**, **stable-privacy**, or **random**.

16.3.11.15. Optional: Configuring host network interfaces for dual port NIC



IMPORTANT

Support for Day 1 operations associated with enabling NIC partitioning for SR-IOV devices is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

Before installation, you can set the **networkConfig** configuration setting in the `install-config.yaml` file to configure host network interfaces using NMState to support dual port NIC.

Prerequisites

- Configure a **PTR** DNS record with a valid hostname for each node with a static IP address.
- Install the NMState CLI (**nmstate**).



NOTE

Errors in the YAML syntax might result in a failure to apply the network configuration. Additionally, maintaining the validated YAML syntax is useful when applying changes using Kubernetes NMState after deployment or when expanding the cluster.

Procedure

1. Add the NMState configuration to the **networkConfig** field to hosts within the **install-config.yaml** file:

```

hosts:
- name: worker-0
  role: worker
  bmc:
    address: redfish+http://<out_of_band_ip>/redfish/v1/Systems/
    username: <user>
    password: <password>
    disableCertificateVerification: false
  bootMACAddress: <NIC1_mac_address>
  bootMode: UEFI
  networkConfig: 1
  interfaces: 2
  - name: eno1 3
    type: ethernet 4
    state: up
    mac-address: 0c:42:a1:55:f3:06
    ipv4:
      enabled: true
      dhcp: false 5
    ethernet:
      sr-iov:
        total-vfs: 2 6
    ipv6:
      enabled: false
      dhcp: false
  - name: sriov:eno1:0
    type: ethernet
    state: up 7
    ipv4:
      enabled: false 8
    ipv6:
      enabled: false
  - name: sriov:eno1:1
    type: ethernet
    state: down
  - name: eno2
    type: ethernet
    state: up

```

```

mac-address: 0c:42:a1:55:f3:07
ipv4:
  enabled: true
ethernet:
  sr-iovs:
    total-vfs: 2
ipv6:
  enabled: false
- name: sriov:eno2:0
  type: ethernet
  state: up
  ipv4:
    enabled: false
  ipv6:
    enabled: false
- name: sriov:eno2:1
  type: ethernet
  state: down
- name: bond0
  type: bond
  state: up
  min-tx-rate: 100 9
  max-tx-rate: 200 10
  link-aggregation:
    mode: active-backup 11
    options:
      primary: sriov:eno1:0 12
    port:
      - sriov:eno1:0
      - sriov:eno2:0
  ipv4:
    address:
      - ip: 10.19.16.57 13
      prefix-length: 23
    dhcp: false
    enabled: true
  ipv6:
    enabled: false
  dns-resolver:
    config:
      server:
        - 10.11.5.160
        - 10.2.70.215
  routes:
    config:
      - destination: 0.0.0.0/0
        next-hop-address: 10.19.17.254
        next-hop-interface: bond0 14
      table-id: 254

```

- 1** The **networkConfig** field contains information about the network configuration of the host, with subfields including **interfaces**, **dns-resolver**, and **routes**.
- 2** The **interfaces** field is an array of network interfaces defined for the host.

- 3 The name of the interface.
- 4 The type of interface. This example creates a ethernet interface.
- 5 Set this to `false` to disable DHCP for the physical function (PF) if it is not strictly required.
- 6 Set to the number of SR-IOV virtual functions (VFs) to instantiate.
- 7 Set this to **up**.
- 8 Set this to **false** to disable IPv4 addressing for the VF attached to the bond.
- 9 Sets a minimum transmission rate, in Mbps, for the VF. This sample value sets a rate of 100 Mbps.
 - This value must be less than or equal to the maximum transmission rate.
 - Intel NICs do not support the **min-tx-rate** parameter. For more information, see [BZ#1772847](#).
- 10 Sets a maximum transmission rate, in Mbps, for the VF. This sample value sets a rate of 200 Mbps.
- 11 Sets the desired bond mode.
- 12 Sets the preferred port of the bonding interface. The primary device is the first of the bonding interfaces to be used and is not abandoned unless it fails. This setting is particularly useful when one NIC in the bonding interface is faster and, therefore, able to handle a bigger load. This setting is only valid when the bonding interface is in active-backup mode (mode 1) and balance-tlb (mode 5).
- 13 Sets a static IP address for the bond interface. This is the node IP address.
- 14 Sets **bond0** as the gateway for the default route.



IMPORTANT

After deploying the cluster, you cannot modify the **networkConfig** configuration setting of **install-config.yaml** file to make changes to the host network interface. Use the Kubernetes NMState Operator to make changes to the host network interface after deployment.

Additional resources

- [Configuring network bonding](#)

16.3.11.16. Configuring multiple cluster nodes

You can simultaneously configure OpenShift Container Platform cluster nodes with identical settings. Configuring multiple cluster nodes avoids adding redundant information for each node to the **install-config.yaml** file. This file contains specific parameters to apply an identical configuration to multiple nodes in the cluster.

Compute nodes are configured separately from the controller node. However, configurations for both node types use the highlighted parameters in the `install-config.yaml` file to enable multi-node configuration. Set the `networkConfig` parameters to **BOND**, as shown in the following example:

```
hosts:
- name: ostest-master-0
[...]
networkConfig: &BOND
interfaces:
- name: bond0
  type: bond
  state: up
  ipv4:
    dhcp: true
    enabled: true
  link-aggregation:
    mode: active-backup
    port:
    - enp2s0
    - enp3s0
- name: ostest-master-1
[...]
networkConfig: *BOND
- name: ostest-master-2
[...]
networkConfig: *BOND
```



NOTE

Configuration of multiple cluster nodes is only available for initial deployments on installer-provisioned infrastructure.

16.3.11.17. Optional: Configuring managed Secure Boot

You can enable managed Secure Boot when deploying an installer-provisioned cluster using Redfish BMC addressing, such as **redfish**, **redfish-virtualmedia**, or **idrac-virtualmedia**. To enable managed Secure Boot, add the `bootMode` configuration setting to each node:

Example

```
hosts:
- name: openshift-master-0
  role: master
  bmc:
    address: redfish://<out_of_band_ip> 1
    username: <username>
    password: <password>
    bootMACAddress: <NIC1_mac_address>
    rootDeviceHints:
      deviceName: "/dev/sda"
    bootMode: UEFISecureBoot 2
```

- 1 Ensure the **bmc.address** setting uses **redfish**, **redfish-virtualmedia**, or **idrac-virtualmedia** as the protocol. See "BMC addressing for HPE iLO" or "BMC addressing for Dell iDRAC" for additional details.
- 2 The **bootMode** setting is **UEFI** by default. Change it to **UEFISecureBoot** to enable managed Secure Boot.

**NOTE**

See "Configuring nodes" in the "Prerequisites" to ensure the nodes can support managed Secure Boot. If the nodes do not support managed Secure Boot, see "Configuring nodes for Secure Boot manually" in the "Configuring nodes" section. Configuring Secure Boot manually requires Redfish virtual media.

**NOTE**

Red Hat does not support Secure Boot with IPMI, because IPMI does not provide Secure Boot management facilities.

16.3.12. Manifest configuration files

16.3.12.1. Creating the OpenShift Container Platform manifests

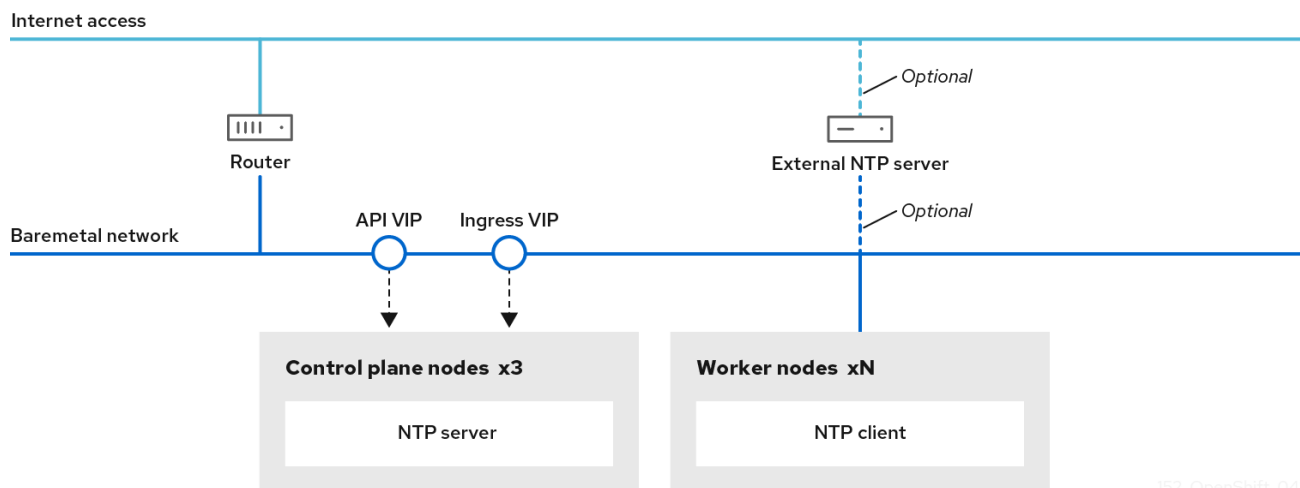
1. Create the OpenShift Container Platform manifests.

```
$ ./openshift-baremetal-install --dir ~/clusterconfigs create manifests
```

```
INFO Consuming Install Config from target directory  
WARNING Making control-plane schedulable by setting MastersSchedulable to true for  
Scheduler cluster settings  
WARNING Discarding the OpenShift Manifest that was provided in the target directory  
because its dependencies are dirty and it needs to be regenerated
```

16.3.12.2. Optional: Configuring NTP for disconnected clusters

OpenShift Container Platform installs the **chrony** Network Time Protocol (NTP) service on the cluster nodes.



152_OpenShift_0421

OpenShift Container Platform nodes must agree on a date and time to run properly. When compute nodes retrieve the date and time from the NTP servers on the control plane nodes, it enables the installation and operation of clusters that are not connected to a routable network and thereby do not have access to a higher stratum NTP server.

Procedure

1. Create a Butane config, **99-master-chrony-conf-override.bu**, including the contents of the **chrony.conf** file for the control plane nodes.



NOTE

See "Creating machine configs with Butane" for information about Butane.

Butane config example

```
variant: openshift
version: 4.16.0
metadata:
  name: 99-master-chrony-conf-override
  labels:
    machineconfiguration.openshift.io/role: master
storage:
files:
  - path: /etc/chrony.conf
    mode: 0644
    overwrite: true
    contents:
      inline: |
        # Use public servers from the pool.ntp.org project.
        # Please consider joining the pool (https://www.pool.ntp.org/join.html).

        # The Machine Config Operator manages this file
        server openshift-master-0.<cluster-name>.<domain> iburst 1
        server openshift-master-1.<cluster-name>.<domain> iburst
        server openshift-master-2.<cluster-name>.<domain> iburst

stratumweight 0
```

```

driftfile /var/lib/chrony/drift
rtcsync
makestep 10 3
bindcmdaddress 127.0.0.1
bindcmdaddress ::1
keyfile /etc/chrony.keys
commandkey 1
generatecommandkey
noclientlog
logchange 0.5
logdir /var/log/chrony

# Configure the control plane nodes to serve as local NTP servers
# for all compute nodes, even if they are not in sync with an
# upstream NTP server.

# Allow NTP client access from the local network.
allow all
# Serve time even if not synchronized to a time source.
local stratum 3 orphan

```

- 1 You must replace **<cluster-name>** with the name of the cluster and replace **<domain>** with the fully qualified domain name.
2. Use Butane to generate a **MachineConfig** object file, **99-master-chrony-conf-override.yaml**, containing the configuration to be delivered to the control plane nodes:

```
$ butane 99-master-chrony-conf-override.bu -o 99-master-chrony-conf-override.yaml
```

3. Create a Butane config, **99-worker-chrony-conf-override.bu**, including the contents of the **chrony.conf** file for the compute nodes that references the NTP servers on the control plane nodes.

Butane config example

```

variant: openshift
version: 4.16.0
metadata:
  name: 99-worker-chrony-conf-override
  labels:
    machineconfiguration.openshift.io/role: worker
storage:
  files:
    - path: /etc/chrony.conf
      mode: 0644
      overwrite: true
      contents:
        inline: |
          # The Machine Config Operator manages this file.
          server openshift-master-0.<cluster-name>.<domain> iburst 1
          server openshift-master-1.<cluster-name>.<domain> iburst
          server openshift-master-2.<cluster-name>.<domain> iburst

          stratumweight 0

```



```
driftfile /var/lib/chrony/drift
rtcsync
makestep 10 3
bindcmdaddress 127.0.0.1
bindcmdaddress ::1
keyfile /etc/chrony.keys
commandkey 1
generatecommandkey
noclientlog
logchange 0.5
logdir /var/log/chrony
```

- 1 You must replace **<cluster-name>** with the name of the cluster and replace **<domain>** with the fully qualified domain name.

4. Use Butane to generate a **MachineConfig** object file, **99-worker-chrony-conf-override.yaml**, containing the configuration to be delivered to the worker nodes:

```
$ butane 99-worker-chrony-conf-override.bu -o 99-worker-chrony-conf-override.yaml
```


16.3.12.3. Configuring network components to run on the control plane

You can configure networking components to run exclusively on the control plane nodes. By default, OpenShift Container Platform allows any node in the machine config pool to host the **ingressVIP** virtual IP address. However, some environments deploy compute nodes in separate subnets from the control plane nodes, which requires configuring the **ingressVIP** virtual IP address to run on the control plane nodes.



IMPORTANT

When deploying remote nodes in separate subnets, you must place the **ingressVIP** virtual IP address exclusively with the control plane nodes.



Procedure

1. Change to the directory storing the **install-config.yaml** file:

```
$ cd ~/clusterconfigs
```

2. Switch to the **manifests** subdirectory:

```
$ cd manifests
```

3. Create a file named **cluster-network-avoid-workers-99-config.yaml**:

```
$ touch cluster-network-avoid-workers-99-config.yaml
```

4. Open the **cluster-network-avoid-workers-99-config.yaml** file in an editor and enter a custom resource (CR) that describes the Operator configuration:

```
apiVersion: machineconfiguration.openshift.io/v1
kind: MachineConfig
metadata:
  name: 50-worker-fix-ipi-rwn
  labels:
    machineconfiguration.openshift.io/role: worker
spec:
  config:
    ignition:
      version: 3.2.0
    storage:
      files:
        - path: /etc/kubernetes/manifests/keepalived.yaml
```

```
mode: 0644
contents:
  source: data,;
```

This manifest places the **ingressVIP** virtual IP address on the control plane nodes. Additionally, this manifest deploys the following processes on the control plane nodes only:

- **openshift-ingress-operator**
- **keepalived**

5. Save the **cluster-network-avoid-workers-99-config.yaml** file.
6. Create a **manifests/cluster-ingress-default-ingresscontroller.yaml** file:

```
apiVersion: operator.openshift.io/v1
kind: IngressController
metadata:
  name: default
  namespace: openshift-ingress-operator
spec:
  nodePlacement:
  nodeSelector:
    matchLabels:
      node-role.kubernetes.io/master: ""
```

7. Consider backing up the **manifests** directory. The installer deletes the **manifests/** directory when creating the cluster.
8. Modify the **cluster-scheduler-02-config.yml** manifest to make the control plane nodes schedulable by setting the **mastersSchedulable** field to **true**. Control plane nodes are not schedulable by default. For example:

```
$ sed -i "s;mastersSchedulable: false;mastersSchedulable: true;g"
clusterconfigs/manifests/cluster-scheduler-02-config.yml
```



NOTE

If control plane nodes are not schedulable after completing this procedure, deploying the cluster will fail.

16.3.12.4. Optional: Deploying routers on compute nodes

During installation, the installation program deploys router pods on compute nodes. By default, the installation program installs two router pods. If a deployed cluster requires additional routers to handle external traffic loads destined for services within the OpenShift Container Platform cluster, you can create a **yaml** file to set an appropriate number of router replicas.



IMPORTANT

Deploying a cluster with only one compute node is not supported. While modifying the router replicas will address issues with the **degraded** state when deploying with one compute node, the cluster loses high availability for the ingress API, which is not suitable for production environments.

**NOTE**

By default, the installation program deploys two routers. If the cluster has no compute nodes, the installation program deploys the two routers on the control plane nodes by default.

Procedure

1. Create a **router-replicas.yaml** file:

```
apiVersion: operator.openshift.io/v1
kind: IngressController
metadata:
  name: default
  namespace: openshift-ingress-operator
spec:
  replicas: <num-of-router-pods>
  endpointPublishingStrategy:
    type: HostNetwork
  nodePlacement:
    nodeSelector:
      matchLabels:
        node-role.kubernetes.io/worker: ""
```

**NOTE**

Replace **<num-of-router-pods>** with an appropriate value. If working with just one compute node, set **replicas**: to **1**. If working with more than 3 compute nodes, you can increase **replicas**: from the default value **2** as appropriate.

2. Save and copy the **router-replicas.yaml** file to the **clusterconfigs/openshift** directory:

```
$ cp ~/router-replicas.yaml clusterconfigs/openshift/99_router-replicas.yaml
```

16.3.12.5. Optional: Configuring the BIOS

The following procedure configures the BIOS during the installation process.


Procedure

1. Create the manifests.
2. Modify the **BareMetalHost** resource file corresponding to the node:

```
$ vim clusterconfigs/openshift/99_openshift-cluster-api_hosts-*.yaml
```

3. Add the BIOS configuration to the **spec** section of the **BareMetalHost** resource:

```
spec:
  firmware:
    simultaneousMultithreadingEnabled: true
    sriovEnabled: true
    virtualizationEnabled: true
```

**NOTE**

Red Hat supports three BIOS configurations. Only servers with BMC type **irmc** are supported. Other types of servers are currently not supported.

4. Create the cluster.

Additional resources

- [Bare metal configuration](#)

16.3.12.6. Optional: Configuring the RAID

The following procedure configures a redundant array of independent disks (RAID) using baseboard management controllers (BMCs) during the installation process.

**NOTE**

If you want to configure a hardware RAID for the node, verify that the node has a supported RAID controller. OpenShift Container Platform 4.16 does not support software RAID.

Table 16.11. Hardware RAID support by vendor

Vendor	BMC and protocol	Firmware version	RAID levels
Fujitsu	iRMC	N/A	0, 1, 5, 6, and 10
Dell	iDRAC with Redfish	Version 6.10.30.20 or later	0, 1, and 5

Procedure

1. Create the manifests.
2. Modify the **BareMetalHost** resource corresponding to the node:

```
$ vim clusterconfigs/openshift/99_openshift-cluster-api_hosts-*.yaml
```

**NOTE**

The following example uses a hardware RAID configuration because OpenShift Container Platform 4.16 does not support software RAID.

- a. If you added a specific RAID configuration to the **spec** section, this causes the node to delete the original RAID configuration in the **preparing** phase and perform a specified configuration on the RAID. For example:

```
spec:
  raid:
    hardwareRAIDVolumes:
```

```
- level: "0" 1
  name: "sda"
  numberOfPhysicalDisks: 1
  rotational: true
  sizeGibibytes: 0
```

1 **level** is a required field, and the others are optional fields.

- b. If you added an empty RAID configuration to the **spec** section, the empty configuration causes the node to delete the original RAID configuration during the **preparing** phase, but does not perform a new configuration. For example:

```
spec:
  raid:
    hardwareRAIDVolumes: []
```

- c. If you do not add a **raid** field in the **spec** section, the original RAID configuration is not deleted, and no new configuration will be performed.

3. Create the cluster.

16.3.12.7. Optional: Configuring storage on nodes

You can make changes to operating systems on OpenShift Container Platform nodes by creating **MachineConfig** objects that are managed by the Machine Config Operator (MCO).

The **MachineConfig** specification includes an ignition config for configuring the machines at first boot. This config object can be used to modify files, systemd services, and other operating system features running on OpenShift Container Platform machines.

Procedure

Use the ignition config to configure storage on nodes. The following **MachineSet** manifest example demonstrates how to add a partition to a device on a primary node. In this example, apply the manifest before installation to have a partition named **recovery** with a size of 16 GiB on the primary node.

1. Create a **custom-partitions.yaml** file and include a **MachineConfig** object that contains your partition layout:

```
apiVersion: machineconfiguration.openshift.io/v1
kind: MachineConfig
metadata:
  labels:
    machineconfiguration.openshift.io/role: primary
  name: 10_primary_storage_config
spec:
  config:
    ignition:
      version: 3.2.0
    storage:
      disks:
        - device: </dev>{x}y{N}
          partitions:
            - label: recovery
              startMiB: 32768
```

```

    sizeMiB: 16384
  filesystems:
  - device: /dev/disk/by-partlabel/recovery
    label: recovery
    format: xfs

```

2. Save and copy the **custom-partitions.yaml** file to the **clusterconfigs/openshift** directory:

```
$ cp ~/<MachineConfig_manifest> ~/clusterconfigs/openshift
```

Additional resources

- [Bare metal configuration](#)
- [Partition naming scheme](#)

16.3.13. Creating a disconnected registry

In some cases, you might want to install an OpenShift Container Platform cluster using a local copy of the installation registry. This could be for enhancing network efficiency because the cluster nodes are on a network that does not have access to the internet.

A local, or mirrored, copy of the registry requires the following:

- A certificate for the registry node. This can be a self-signed certificate.
- A web server that a container on a system will serve.
- An updated pull secret that contains the certificate and local repository information.



NOTE

Creating a disconnected registry on a registry node is optional. If you need to create a disconnected registry on a registry node, you must complete all of the following subsections.

Prerequisites

- If you have already prepared a mirror registry for [Mirroring images for a disconnected installation](#), you can skip directly to [Modify the install-config.yaml file to use the disconnected registry](#).

16.3.13.1. Preparing the registry node to host the mirrored registry

The following steps must be completed prior to hosting a mirrored registry on bare metal.

Procedure

1. Open the firewall port on the registry node:

```
$ sudo firewall-cmd --add-port=5000/tcp --zone=libvirt --permanent
```

```
$ sudo firewall-cmd --add-port=5000/tcp --zone=public --permanent
```

```
$ sudo firewall-cmd --reload
```

2. Install the required packages for the registry node:

```
$ sudo yum -y install python3 podman httpd httpd-tools jq
```

3. Create the directory structure where the repository information will be held:

```
$ sudo mkdir -p /opt/registry/{auth,certs,data}
```

16.3.13.2. Mirroring the OpenShift Container Platform image repository for a disconnected registry

Complete the following steps to mirror the OpenShift Container Platform image repository for a disconnected registry.

Prerequisites

- Your mirror host has access to the internet.
- You configured a mirror registry to use in your restricted network and can access the certificate and credentials that you configured.
- You downloaded the [pull secret from Red Hat OpenShift Cluster Manager](#) and modified it to include authentication to your mirror repository.

Procedure

1. Review the [OpenShift Container Platform downloads page](#) to determine the version of OpenShift Container Platform that you want to install and determine the corresponding tag on the [Repository Tags](#) page.
2. Set the required environment variables:
 - a. Export the release version:

```
$ OCP_RELEASE=<release_version>
```

For **<release_version>**, specify the tag that corresponds to the version of OpenShift Container Platform to install, such as **4.5.4**.

- b. Export the local registry name and host port:

```
$ LOCAL_REGISTRY='<local_registry_host_name>:<local_registry_host_port>'
```

For **<local_registry_host_name>**, specify the registry domain name for your mirror repository, and for **<local_registry_host_port>**, specify the port that it serves content on.

- c. Export the local repository name:

```
$ LOCAL_REPOSITORY='<local_repository_name>'
```


For **<local_repository_name>**, specify the name of the repository to create in your registry, such as **ocp4/openshift4**.

- d. Export the name of the repository to mirror:

```
$ PRODUCT_REPO='openshift-release-dev'
```

For a production release, you must specify **openshift-release-dev**.

- e. Export the path to your registry pull secret:

```
$ LOCAL_SECRET_JSON='<path_to_pull_secret>'
```

For **<path_to_pull_secret>**, specify the absolute path to and file name of the pull secret for your mirror registry that you created.

- f. Export the release mirror:

```
$ RELEASE_NAME="ocp-release"
```

For a production release, you must specify **ocp-release**.

- g. Export the type of architecture for your cluster:

```
$ ARCHITECTURE=<cluster_architecture> 1
```

- 1 Specify the architecture of the cluster, such as **x86_64**, **aarch64**, **s390x**, or **ppc64le**.

- h. Export the path to the directory to host the mirrored images:

```
$ REMOVABLE_MEDIA_PATH=<path> 1
```

- 1 Specify the full path, including the initial forward slash (/) character.

3. Mirror the version images to the mirror registry:

- If your mirror host does not have internet access, take the following actions:
 - i. Connect the removable media to a system that is connected to the internet.
 - ii. Review the images and configuration manifests to mirror:

```
$ oc adm release mirror -a ${LOCAL_SECRET_JSON} \
  --from=quay.io/${PRODUCT_REPO}/${RELEASE_NAME}:${OCP_RELEASE}-
  ${ARCHITECTURE} \
  --to=${LOCAL_REGISTRY}/${LOCAL_REPOSITORY} \
  --to-release-
  image=${LOCAL_REGISTRY}/${LOCAL_REPOSITORY}:${OCP_RELEASE}-
  ${ARCHITECTURE} --dry-run
```

- iii. Record the entire **imageContentSources** section from the output of the previous command. The information about your mirrors is unique to your mirrored repository, and you must add the **imageContentSources** section to the **install-config.yaml** file during

installation.

- iv. Mirror the images to a directory on the removable media:

```
$ oc adm release mirror -a ${LOCAL_SECRET_JSON} --to-dir=${REMOVABLE_MEDIA_PATH}/mirror quay.io/${PRODUCT_REPO}/${RELEASE_NAME}:${OCP_RELEASE}-${ARCHITECTURE}
```

- v. Take the media to the restricted network environment and upload the images to the local container registry.

```
$ oc image mirror -a ${LOCAL_SECRET_JSON} --from-dir=${REMOVABLE_MEDIA_PATH}/mirror "file://openshift/release:${OCP_RELEASE}*" ${LOCAL_REGISTRY}/${LOCAL_REPOSITORY} 1
```

- 1** For **REMOVABLE_MEDIA_PATH**, you must use the same path that you specified when you mirrored the images.

- If the local container registry is connected to the mirror host, take the following actions:
 - i. Directly push the release images to the local registry by using following command:

```
$ oc adm release mirror -a ${LOCAL_SECRET_JSON} \ --from=quay.io/${PRODUCT_REPO}/${RELEASE_NAME}:${OCP_RELEASE}-${ARCHITECTURE} \ --to=${LOCAL_REGISTRY}/${LOCAL_REPOSITORY} \ --to-release-image=${LOCAL_REGISTRY}/${LOCAL_REPOSITORY}:${OCP_RELEASE}-${ARCHITECTURE}
```

This command pulls the release information as a digest, and its output includes the **imageContentSources** data that you require when you install your cluster.

- ii. Record the entire **imageContentSources** section from the output of the previous command. The information about your mirrors is unique to your mirrored repository, and you must add the **imageContentSources** section to the **install-config.yaml** file during installation.



NOTE

The image name gets patched to Quay.io during the mirroring process, and the podman images will show Quay.io in the registry on the bootstrap virtual machine.

4. To create the installation program that is based on the content that you mirrored, extract it and pin it to the release:

- If your mirror host does not have internet access, run the following command:

```
$ oc adm release extract -a ${LOCAL_SECRET_JSON} --command=openshift-baremetal-install "${LOCAL_REGISTRY}/${LOCAL_REPOSITORY}:${OCP_RELEASE}"
```

- If the local container registry is connected to the mirror host, run the following command:

```
$ oc adm release extract -a ${LOCAL_SECRET_JSON} --command=openshift-
baremetal-install "${LOCAL_REGISTRY}/${LOCAL_REPOSITORY}:${OCP_RELEASE}-
${ARCHITECTURE}"
```



IMPORTANT

To ensure that you use the correct images for the version of OpenShift Container Platform that you selected, you must extract the installation program from the mirrored content.

You must perform this step on a machine with an active internet connection.

If you are in a disconnected environment, use the **--image** flag as part of `must-gather` and point to the payload image.

5. For clusters using installer-provisioned infrastructure, run the following command:

```
$ openshift-baremetal-install
```

16.3.13.3. Modify the `install-config.yaml` file to use the disconnected registry

On the provisioner node, the `install-config.yaml` file should use the newly created pull-secret from the `pull-secret-update.txt` file. The `install-config.yaml` file must also contain the disconnected registry node's certificate and registry information.

Procedure

1. Add the disconnected registry node's certificate to the `install-config.yaml` file:

```
$ echo "additionalTrustBundle: |" >> install-config.yaml
```

The certificate should follow the `"additionalTrustBundle: |"` line and be properly indented, usually by two spaces.

```
$ sed -e 's/^ /' /opt/registry/certs/domain.crt >> install-config.yaml
```

2. Add the mirror information for the registry to the `install-config.yaml` file:

```
$ echo "imageContentSources:" >> install-config.yaml
```

```
$ echo "- mirrors:" >> install-config.yaml
```

```
$ echo "  - registry.example.com:5000/ocp4/openshift4" >> install-config.yaml
```

Replace **registry.example.com** with the registry's fully qualified domain name.

```
$ echo "    source: quay.io/openshift-release-dev/ocp-release" >> install-config.yaml
```

```
$ echo "- mirrors:" >> install-config.yaml
```

```
$ echo " - registry.example.com:5000/ocp4/openshift4" >> install-config.yaml
```

Replace **registry.example.com** with the registry's fully qualified domain name.

```
$ echo " source: quay.io/openshift-release-dev/ocp-v4.0-art-dev" >> install-config.yaml
```

16.3.14. Validation checklist for installation

- OpenShift Container Platform installer has been retrieved.
- OpenShift Container Platform installer has been extracted.
- Required parameters for the **install-config.yaml** have been configured.
- The **hosts** parameter for the **install-config.yaml** has been configured.
- The **bmc** parameter for the **install-config.yaml** has been configured.
- Conventions for the values configured in the **bmc address** field have been applied.
- Created the OpenShift Container Platform manifests.
- (Optional) Deployed routers on compute nodes.
- (Optional) Created a disconnected registry.
- (Optional) Validate disconnected registry settings if in use.

16.3.15. Deploying the cluster via the OpenShift Container Platform installer

Run the OpenShift Container Platform installer:

```
$ ./openshift-baremetal-install --dir ~/clusterconfigs --log-level debug create cluster
```

16.3.16. Following the installation

During the deployment process, you can check the installation's overall status by issuing the **tail** command to the **.openshift_install.log** log file in the install directory folder:

```
$ tail -f /path/to/install-dir/.openshift_install.log
```

16.3.17. Verifying static IP address configuration

If the DHCP reservation for a cluster node specifies an infinite lease, after the installer successfully provisions the node, the dispatcher script checks the node's network configuration. If the script determines that the network configuration contains an infinite DHCP lease, it creates a new connection using the IP address of the DHCP lease as a static IP address.



NOTE

The dispatcher script might run on successfully provisioned nodes while the provisioning of other nodes in the cluster is ongoing.

Verify the network configuration is working properly.

Procedure

1. Check the network interface configuration on the node.
2. Turn off the DHCP server and reboot the OpenShift Container Platform node and ensure that the network configuration works properly.

16.3.18. Preparing to reinstall a cluster on bare metal

Before you reinstall a cluster on bare metal, you must perform cleanup operations.

Procedure

1. Remove or reformat the disks for the bootstrap, control plane node, and compute nodes. If you are working in a hypervisor environment, you must add any disks you removed.
2. Delete the artifacts that the previous installation generated:

```
$ cd ; /bin/rm -rf auth/ bootstrap.ign master.ign worker.ign metadata.json \
.openshift_install.log .openshift_install_state.json
```

3. Generate new manifests and Ignition config files. See "Creating the Kubernetes manifest and Ignition config files" for more information.
4. Upload the new bootstrap, control plane, and compute node Ignition config files that the installation program created to your HTTP server. This will overwrite the previous Ignition files.

16.3.19. Additional resources


- [OpenShift Container Platform Creating the Kubernetes manifest and Ignition config files](#)
- [Understanding update channels and releases](#)

16.4. INSTALLER-PROVISIONED POSTINSTALLATION CONFIGURATION

After successfully deploying an installer-provisioned cluster, consider the following postinstallation procedures.

16.4.1. Optional: Configuring NTP for disconnected clusters

OpenShift Container Platform installs the **chrony** Network Time Protocol (NTP) service on the cluster nodes. Use the following procedure to configure NTP servers on the control plane nodes and configure compute nodes as NTP clients of the control plane nodes after a successful deployment.



152_OpenShift_0421

OpenShift Container Platform nodes must agree on a date and time to run properly. When compute nodes retrieve the date and time from the NTP servers on the control plane nodes, it enables the installation and operation of clusters that are not connected to a routable network and thereby do not have access to a higher stratum NTP server.

Procedure

1. Create a Butane config, **99-master-chrony-conf-override.bu**, including the contents of the **chrony.conf** file for the control plane nodes.



NOTE

See "Creating machine configs with Butane" for information about Butane.

Butane config example

```
variant: openshift
version: 4.16.0
metadata:
  name: 99-master-chrony-conf-override
  labels:
    machineconfiguration.openshift.io/role: master
storage:
files:
  - path: /etc/chrony.conf
    mode: 0644
    overwrite: true
    contents:
      inline: |
        # Use public servers from the pool.ntp.org project.
        # Please consider joining the pool (https://www.pool.ntp.org/join.html).

        # The Machine Config Operator manages this file
        server openshift-master-0.<cluster-name>.<domain> iburst 1
        server openshift-master-1.<cluster-name>.<domain> iburst
        server openshift-master-2.<cluster-name>.<domain> iburst

        stratumweight 0
```

```

driftfile /var/lib/chrony/drift
rtcsync
makestep 10 3
bindcmdaddress 127.0.0.1
bindcmdaddress ::1
keyfile /etc/chrony.keys
commandkey 1
generatecommandkey
noclientlog
logchange 0.5
logdir /var/log/chrony

# Configure the control plane nodes to serve as local NTP servers
# for all compute nodes, even if they are not in sync with an
# upstream NTP server.

# Allow NTP client access from the local network.
allow all
# Serve time even if not synchronized to a time source.
local stratum 3 orphan

```

- 1 You must replace **<cluster-name>** with the name of the cluster and replace **<domain>** with the fully qualified domain name.
2. Use Butane to generate a **MachineConfig** object file, **99-master-chrony-conf-override.yaml**, containing the configuration to be delivered to the control plane nodes:

```
$ butane 99-master-chrony-conf-override.bu -o 99-master-chrony-conf-override.yaml
```

3. Create a Butane config, **99-worker-chrony-conf-override.bu**, including the contents of the **chrony.conf** file for the compute nodes that references the NTP servers on the control plane nodes.

Butane config example

```

variant: openshift
version: 4.16.0
metadata:
  name: 99-worker-chrony-conf-override
  labels:
    machineconfiguration.openshift.io/role: worker
storage:
  files:
    - path: /etc/chrony.conf
      mode: 0644
      overwrite: true
      contents:
        inline: |
          # The Machine Config Operator manages this file.
          server openshift-master-0.<cluster-name>.<domain> iburst 1
          server openshift-master-1.<cluster-name>.<domain> iburst
          server openshift-master-2.<cluster-name>.<domain> iburst

          stratumweight 0

```

```
driftfile /var/lib/chrony/drift
rtcsync
makestep 10 3
bindcmdaddress 127.0.0.1
bindcmdaddress ::1
keyfile /etc/chrony.keys
commandkey 1
generatecommandkey
noclientlog
logchange 0.5
logdir /var/log/chrony
```

- 1 You must replace **<cluster-name>** with the name of the cluster and replace **<domain>** with the fully qualified domain name.

4. Use Butane to generate a **MachineConfig** object file, **99-worker-chrony-conf-override.yaml**, containing the configuration to be delivered to the worker nodes:

```
$ butane 99-worker-chrony-conf-override.bu -o 99-worker-chrony-conf-override.yaml
```

5. Apply the **99-master-chrony-conf-override.yaml** policy to the control plane nodes.

```
$ oc apply -f 99-master-chrony-conf-override.yaml
```

Example output

```
machineconfig.machineconfiguration.openshift.io/99-master-chrony-conf-override created
```

6. Apply the **99-worker-chrony-conf-override.yaml** policy to the compute nodes.

```
$ oc apply -f 99-worker-chrony-conf-override.yaml
```

Example output

```
machineconfig.machineconfiguration.openshift.io/99-worker-chrony-conf-override created
```

7. Check the status of the applied NTP settings.

```
$ oc describe machineconfigpool
```

16.4.2. Enabling a provisioning network after installation

The assisted installer and installer-provisioned installation for bare metal clusters provide the ability to deploy a cluster without a **provisioning** network. This capability is for scenarios such as proof-of-concept clusters or deploying exclusively with Redfish virtual media when each node's baseboard management controller is routable via the **baremetal** network.

You can enable a **provisioning** network after installation using the Cluster Baremetal Operator (CBO).

Prerequisites

- A dedicated physical network must exist, connected to all worker and control plane nodes.
- You must isolate the native, untagged physical network.
- The network cannot have a DHCP server when the **provisioningNetwork** configuration setting is set to **Managed**.
- You can omit the **provisioningInterface** setting in OpenShift Container Platform 4.10 to use the **bootMACAddress** configuration setting.

Procedure

1. When setting the **provisioningInterface** setting, first identify the provisioning interface name for the cluster nodes. For example, **eth0** or **eno1**.
2. Enable the Preboot eXecution Environment (PXE) on the **provisioning** network interface of the cluster nodes.
3. Retrieve the current state of the **provisioning** network and save it to a provisioning custom resource (CR) file:

```
$ oc get provisioning -o yaml > enable-provisioning-nw.yaml
```

4. Modify the provisioning CR file:

```
$ vim ~/enable-provisioning-nw.yaml
```

Scroll down to the **provisioningNetwork** configuration setting and change it from **Disabled** to **Managed**. Then, add the **provisioningIP**, **provisioningNetworkCIDR**, **provisioningDHCPRange**, **provisioningInterface**, and **watchAllNameSpaces** configuration settings after the **provisioningNetwork** setting. Provide appropriate values for each setting.

```
apiVersion: v1
items:
- apiVersion: metal3.io/v1alpha1
  kind: Provisioning
  metadata:
    name: provisioning-configuration
  spec:
    provisioningNetwork: 1
    provisioningIP: 2
    provisioningNetworkCIDR: 3
    provisioningDHCPRange: 4
    provisioningInterface: 5
    watchAllNameSpaces: 6
```

1 The **provisioningNetwork** is one of **Managed**, **Unmanaged**, or **Disabled**. When set to **Managed**, Metal3 manages the provisioning network and the CBO deploys the Metal3 pod with a configured DHCP server. When set to **Unmanaged**, the system administrator configures the DHCP server manually.

2 The **provisioningIP** is the static IP address that the DHCP server and ironic use to provision the network. This static IP address must be within the **provisioning** subnet, and outside of the DHCP range. If you configure this setting, it must have a valid IP address

even if the **provisioning** network is **Disabled**. The static IP address is bound to the metal3 pod. If the metal3 pod fails and moves to another server, the static IP address also moves to the new server.

- 3 The Classless Inter-Domain Routing (CIDR) address. If you configure this setting, it must have a valid CIDR address even if the **provisioning** network is **Disabled**. For example: **192.168.0.1/24**.
- 4 The DHCP range. This setting is only applicable to a **Managed** provisioning network. Omit this configuration setting if the **provisioning** network is **Disabled**. For example: **192.168.0.64, 192.168.0.253**.
- 5 The NIC name for the **provisioning** interface on cluster nodes. The **provisioningInterface** setting is only applicable to **Managed** and **Unmanaged** provisioning networks. Omit the **provisioningInterface** configuration setting if the **provisioning** network is **Disabled**. Omit the **provisioningInterface** configuration setting to use the **bootMACAddress** configuration setting instead.
- 6 Set this setting to **true** if you want metal3 to watch namespaces other than the default **openshift-machine-api** namespace. The default value is **false**.

5. Save the changes to the provisioning CR file.

6. Apply the provisioning CR file to the cluster:

```
$ oc apply -f enable-provisioning-nw.yaml
```

16.5. EXPANDING THE CLUSTER

After deploying an installer-provisioned OpenShift Container Platform cluster, you can use the following procedures to expand the number of worker nodes. Ensure that each prospective worker node meets the prerequisites.



NOTE

Expanding the cluster using RedFish Virtual Media involves meeting minimum firmware requirements. See **Firmware requirements for installing with virtual media** in the **Prerequisites** section for additional details when expanding the cluster using RedFish Virtual Media.

16.5.1. Preparing the bare metal node

To expand your cluster, you must provide the node with the relevant IP address. This can be done with a static configuration, or with a DHCP (Dynamic Host Configuration protocol) server. When expanding the cluster using a DHCP server, each node must have a DHCP reservation.



RESERVING IP ADDRESSES SO THEY BECOME STATIC IP ADDRESSES

Some administrators prefer to use static IP addresses so that each node's IP address remains constant in the absence of a DHCP server. To configure static IP addresses with NMState, see "Optional: Configuring host network interfaces in the **install-config.yaml** file" in the "Setting up the environment for an OpenShift installation" section for additional details.

Preparing the bare metal node requires executing the following procedure from the provisioner node.

Procedure

1. Get the **oc** binary:

```
$ curl -s https://mirror.openshift.com/pub/openshift-v4/clients/ocp/$VERSION/openshift-client-linux-$VERSION.tar.gz | tar zxvf - oc
```

```
$ sudo cp oc /usr/local/bin
```

2. Power off the bare metal node by using the baseboard management controller (BMC), and ensure it is off.
3. Retrieve the user name and password of the bare metal node's baseboard management controller. Then, create **base64** strings from the user name and password:

```
$ echo -ne "root" | base64
```

```
$ echo -ne "password" | base64
```

4. Create a configuration file for the bare metal node. Depending on whether you are using a static configuration or a DHCP server, use one of the following example **bmh.yaml** files, replacing values in the YAML to match your environment:

```
$ vim bmh.yaml
```

- **Static configuration bmh.yaml:**

```
---
apiVersion: v1 1
kind: Secret
metadata:
  name: openshift-worker-<num>-network-config-secret 2
  namespace: openshift-machine-api
type: Opaque
stringData:
  nmstate: | 3
  interfaces: 4
    - name: <nic1_name> 5
      type: ethernet
      state: up
      ipv4:
        address:
          - ip: <ip_address> 6
            prefix-length: 24
            enabled: true
  dns-resolver:
    config:
      server:
        - <dns_ip_address> 7
  routes:
    config:
```

```

- destination: 0.0.0.0/0
  next-hop-address: <next_hop_ip_address> 8
  next-hop-interface: <next_hop_nic1_name> 9
---
apiVersion: v1
kind: Secret
metadata:
  name: openshift-worker-<num>-bmc-secret 10
  namespace: openshift-machine-api
type: Opaque
data:
  username: <base64_of_uid> 11
  password: <base64_of_pwd> 12
---
apiVersion: metal3.io/v1alpha1
kind: BareMetalHost
metadata:
  name: openshift-worker-<num> 13
  namespace: openshift-machine-api
spec:
  online: True
  bootMACAddress: <nic1_mac_address> 14
  bmc:
    address: <protocol>://<bmc_url> 15
    credentialsName: openshift-worker-<num>-bmc-secret 16
    disableCertificateVerification: True 17
    username: <bmc_username> 18
    password: <bmc_password> 19
  rootDeviceHints:
    deviceName: <root_device_hint> 20
  preprovisioningNetworkDataName: openshift-worker-<num>-network-config-secret 21

```

- 1 To configure the network interface for a newly created node, specify the name of the secret that contains the network configuration. Follow the **nmstate** syntax to define the network configuration for your node. See "Optional: Configuring host network interfaces in the install-config.yaml file" for details on configuring NMState syntax.
- 2 10 13 16 Replace **<num>** for the worker number of the bare metal node in the **name** fields, the **credentialsName** field, and the **preprovisioningNetworkDataName** field.
- 3 Add the NMState YAML syntax to configure the host interfaces.
- 4 Optional: If you have configured the network interface with **nmstate**, and you want to disable an interface, set **state: up** with the IP addresses set to **enabled: false** as shown:

```

---
interfaces:
- name: <nic_name>
  type: ethernet
  state: up
  ipv4:

```

```

enabled: false
ipv6:
  enabled: false

```

- 5 6 7 8 9 Replace `<nic1_name>`, `<ip_address>`, `<dns_ip_address>`, `<next_hop_ip_address>` and `<next_hop_nic1_name>` with appropriate values.
- 11 12 Replace `<base64_of_uid>` and `<base64_of_pwd>` with the base64 string of the user name and password.
- 14 Replace `<nic1_mac_address>` with the MAC address of the bare metal node's first NIC. See the "BMC addressing" section for additional BMC configuration options.
- 15 Replace `<protocol>` with the BMC protocol, such as IPMI, RedFish, or others. Replace `<bmc_url>` with the URL of the bare metal node's baseboard management controller.
- 17 To skip certificate validation, set `disableCertificateVerification` to true.
- 18 19 Replace `<bmc_username>` and `<bmc_password>` with the string of the BMC user name and password.
- 20 Optional: Replace `<root_device_hint>` with a device path if you specify a root device hint.
- 21 Optional: If you have configured the network interface for the newly created node, provide the network configuration secret name in the `preprovisioningNetworkDataName` of the BareMetalHost CR.

- DHCP configuration `bmh.yaml`:

```

---
apiVersion: v1
kind: Secret
metadata:
  name: openshift-worker-<num>-bmc-secret 1
  namespace: openshift-machine-api
type: Opaque
data:
  username: <base64_of_uid> 2
  password: <base64_of_pwd> 3
---
apiVersion: metal3.io/v1alpha1
kind: BareMetalHost
metadata:
  name: openshift-worker-<num> 4
  namespace: openshift-machine-api
spec:
  online: True
  bootMACAddress: <nic1_mac_address> 5
  bmc:
    address: <protocol>://<bmc_url> 6
    credentialsName: openshift-worker-<num>-bmc-secret 7
    disableCertificateVerification: True 8

```

```

username: <bmc_username> 9
password: <bmc_password> 10
rootDeviceHints:
  deviceName: <root_device_hint> 11
preprovisioningNetworkDataName: openshift-worker-<num>-network-config-secret 12

```

- 1** **4** **7** Replace **<num>** for the worker number of the bare metal node in the **name** fields, the **credentialsName** field, and the **preprovisioningNetworkDataName** field.
- 2** **3** Replace **<base64_of_uid>** and **<base64_of_pwd>** with the base64 string of the user name and password.
- 5** Replace **<nic1_mac_address>** with the MAC address of the bare metal node's first NIC. See the "BMC addressing" section for additional BMC configuration options.
- 6** Replace **<protocol>** with the BMC protocol, such as IPMI, RedFish, or others. Replace **<bmc_url>** with the URL of the bare metal node's baseboard management controller.
- 8** To skip certificate validation, set **disableCertificateVerification** to true.
- 9** **10** Replace **<bmc_username>** and **<bmc_password>** with the string of the BMC user name and password.
- 11** Optional: Replace **<root_device_hint>** with a device path if you specify a root device hint.
- 12** Optional: If you have configured the network interface for the newly created node, provide the network configuration secret name in the **preprovisioningNetworkDataName** of the BareMetalHost CR.



NOTE

If the MAC address of an existing bare metal node matches the MAC address of a bare metal host that you are attempting to provision, then the Ironic installation will fail. If the host enrollment, inspection, cleaning, or other Ironic steps fail, the Bare Metal Operator retries the installation continuously. See "Diagnosing a host duplicate MAC address" for more information.

5. Create the bare metal node:

```
$ oc -n openshift-machine-api create -f bmh.yaml
```

Example output

```
secret/openshift-worker-<num>-network-config-secret created
secret/openshift-worker-<num>-bmc-secret created
baremetalhost.metal3.io/openshift-worker-<num> created
```

Where **<num>** will be the worker number.

6. Power up and inspect the bare metal node:

```
$ oc -n openshift-machine-api get bmh openshift-worker-<num>
```

Where **<num>** is the worker node number.

Example output

```
NAME                STATE    CONSUMER ONLINE  ERROR
openshift-worker-<num> available      true
```



NOTE

To allow the worker node to join the cluster, scale the **machineset** object to the number of the **BareMetalHost** objects. You can scale nodes either manually or automatically. To scale nodes automatically, use the **metal3.io/autoscale-to-hosts** annotation for **machineset**.

Additional resources

- See [Optional: Configuring host network interfaces in the install-config.yaml file](#) for details on configuring the NMState syntax.
- See [Automatically scaling machines to the number of available bare metal hosts](#) for details on automatically scaling machines.

16.5.2. Replacing a bare-metal control plane node

Use the following procedure to replace an installer-provisioned OpenShift Container Platform control plane node.



IMPORTANT

If you reuse the **BareMetalHost** object definition from an existing control plane host, do not leave the **externallyProvisioned** field set to **true**.

Existing control plane **BareMetalHost** objects may have the **externallyProvisioned** flag set to **true** if they were provisioned by the OpenShift Container Platform installation program.

Prerequisites

- You have access to the cluster as a user with the **cluster-admin** role.
- You have taken an etcd backup.



IMPORTANT

Take an etcd backup before performing this procedure so that you can restore your cluster if you encounter any issues. For more information about taking an etcd backup, see the *Additional resources* section.

Procedure

1. Ensure that the Bare Metal Operator is available:

```
$ oc get clusteroperator baremetal
```

Example output

```

NAME      VERSION AVAILABLE PROGRESSING DEGRADED SINCE MESSAGE
baremetal 4.16 True      False      False      3d15h

```

- Remove the old **BareMetalHost** and **Machine** objects:

```

$ oc delete bmh -n openshift-machine-api <host_name>
$ oc delete machine -n openshift-machine-api <machine_name>

```

Replace **<host_name>** with the name of the host and **<machine_name>** with the name of the machine. The machine name appears under the **CONSUMER** field.

After you remove the **BareMetalHost** and **Machine** objects, then the machine controller automatically deletes the **Node** object.

- Create the new **BareMetalHost** object and the secret to store the BMC credentials:

```

$ cat <<EOF | oc apply -f -
apiVersion: v1
kind: Secret
metadata:
  name: control-plane-<num>-bmc-secret 1
  namespace: openshift-machine-api
data:
  username: <base64_of_uid> 2
  password: <base64_of_pwd> 3
type: Opaque
---
apiVersion: metal3.io/v1alpha1
kind: BareMetalHost
metadata:
  name: control-plane-<num> 4
  namespace: openshift-machine-api
spec:
  automatedCleaningMode: disabled
  bmc:
    address: <protocol>://<bmc_ip> 5
    credentialsName: control-plane-<num>-bmc-secret 6
  bootMACAddress: <NIC1_mac_address> 7
  bootMode: UEFI
  externallyProvisioned: false
  online: true
EOF

```

1 4 6 Replace **<num>** for the control plane number of the bare metal node in the **name** fields and the **credentialsName** field.

2 Replace **<base64_of_uid>** with the **base64** string of the user name.

3 Replace **<base64_of_pwd>** with the **base64** string of the password.

5

Replace **<protocol>** with the BMC protocol, such as **redfish**, **redfish-virtualmedia**, **idrac-virtualmedia**, or others. Replace **<bmc_ip>** with the IP address of the bare metal node's

- 7 Replace **<NIC1_mac_address>** with the MAC address of the bare metal node's first NIC.

After the inspection is complete, the **BareMetalHost** object is created and available to be provisioned.

4. View available **BareMetalHost** objects:

```
$ oc get bmh -n openshift-machine-api
```

Example output

NAME	STATE	CONSUMER	ONLINE	ERROR	AGE
control-plane-1.example.com	available	control-plane-1	true		1h10m
control-plane-2.example.com	externally provisioned	control-plane-2		true	4h53m
control-plane-3.example.com	externally provisioned	control-plane-3		true	4h53m
compute-1.example.com	provisioned	compute-1-ktmmx	true		4h53m
compute-1.example.com	provisioned	compute-2-l2zmb	true		4h53m

There are no **MachineSet** objects for control plane nodes, so you must create a **Machine** object instead. You can copy the **providerSpec** from another control plane **Machine** object.

5. Create a **Machine** object:

```
$ cat <<EOF | oc apply -f -
apiVersion: machine.openshift.io/v1beta1
kind: Machine
metadata:
  annotations:
    metal3.io/BareMetalHost: openshift-machine-api/control-plane-<num> 1
  labels:
    machine.openshift.io/cluster-api-cluster: control-plane-<num> 2
    machine.openshift.io/cluster-api-machine-role: master
    machine.openshift.io/cluster-api-machine-type: master
  name: control-plane-<num> 3
  namespace: openshift-machine-api
spec:
  metadata: {}
  providerSpec:
    value:
      apiVersion: baremetal.cluster.k8s.io/v1alpha1
      customDeploy:
        method: install_coreos
      hostSelector: {}
      image:
        checksum: ""
        url: ""
      kind: BareMetalMachineProviderSpec
```

```

metadata:
  creationTimestamp: null
  userData:
    name: master-user-data-managed
EOF

```

- 1
 - 2
 - 3
- Replace **<num>** for the control plane number of the bare metal node in the **name**, **labels** and **annotations** fields.

6. To view the **BareMetalHost** objects, run the following command:

```
$ oc get bmh -A
```

Example output

NAME	STATE	CONSUMER	ONLINE	ERROR	AGE
control-plane-1.example.com	provisioned	control-plane-1	true		2h53m
control-plane-2.example.com	externally provisioned	control-plane-2		true	5h53m
control-plane-3.example.com	externally provisioned	control-plane-3		true	5h53m
compute-1.example.com	provisioned	compute-1-ktmmx		true	5h53m
compute-2.example.com	provisioned	compute-2-l2zmb		true	5h53m

7. After the RHCOS installation, verify that the **BareMetalHost** is added to the cluster:

```
$ oc get nodes
```

Example output

NAME	STATUS	ROLES	AGE	VERSION
control-plane-1.example.com	available	master	4m2s	v1.29.4
control-plane-2.example.com	available	master	141m	v1.29.4
control-plane-3.example.com	available	master	141m	v1.29.4
compute-1.example.com	available	worker	87m	v1.29.4
compute-2.example.com	available	worker	87m	v1.29.4



NOTE

After replacement of the new control plane node, the etcd pod running in the new node is in **crashloopback** status. See "Replacing an unhealthy etcd member" in the *Additional resources* section for more information.

Additional resources

- [Replacing an unhealthy etcd member](#)
- [Backing up etcd](#)
- [Bare metal configuration](#)

- [BMC addressing](#)

16.5.3. Preparing to deploy with Virtual Media on the baremetal network

If the **provisioning** network is enabled and you want to expand the cluster using Virtual Media on the **baremetal** network, use the following procedure.

Prerequisites

- There is an existing cluster with a **baremetal** network and a **provisioning** network.

Procedure

1. Edit the **provisioning** custom resource (CR) to enable deploying with Virtual Media on the **baremetal** network:

```
oc edit provisioning

apiVersion: metal3.io/v1alpha1
kind: Provisioning
metadata:
  creationTimestamp: "2021-08-05T18:51:50Z"
  finalizers:
  - provisioning.metal3.io
  generation: 8
  name: provisioning-configuration
  resourceVersion: "551591"
  uid: f76e956f-24c6-4361-aa5b-feaf72c5b526
spec:
  provisioningDHCPRange: 172.22.0.10,172.22.0.254
  provisioningIP: 172.22.0.3
  provisioningInterface: enp1s0
  provisioningNetwork: Managed
  provisioningNetworkCIDR: 172.22.0.0/24
  virtualMediaViaExternalNetwork: true 1
status:
  generations:
  - group: apps
    hash: ""
    lastGeneration: 7
    name: metal3
    namespace: openshift-machine-api
    resource: deployments
  - group: apps
    hash: ""
    lastGeneration: 1
    name: metal3-image-cache
    namespace: openshift-machine-api
    resource: daemonsets
  observedGeneration: 8
  readyReplicas: 0
```

- 1** Add **virtualMediaViaExternalNetwork: true** to the **provisioning** CR.

- If the image URL exists, edit the **machineset** to use the API VIP address. This step only applies to clusters installed in versions 4.9 or earlier.

```
oc edit machineset
```

```
apiVersion: machine.openshift.io/v1beta1
kind: MachineSet
metadata:
  creationTimestamp: "2021-08-05T18:51:52Z"
  generation: 11
  labels:
    machine.openshift.io/cluster-api-cluster: ostest-hwmdt
    machine.openshift.io/cluster-api-machine-role: worker
    machine.openshift.io/cluster-api-machine-type: worker
  name: ostest-hwmdt-worker-0
  namespace: openshift-machine-api
  resourceVersion: "551513"
  uid: fad1c6e0-b9da-4d4a-8d73-286f78788931
spec:
  replicas: 2
  selector:
    matchLabels:
      machine.openshift.io/cluster-api-cluster: ostest-hwmdt
      machine.openshift.io/cluster-api-machineset: ostest-hwmdt-worker-0
  template:
    metadata:
      labels:
        machine.openshift.io/cluster-api-cluster: ostest-hwmdt
        machine.openshift.io/cluster-api-machine-role: worker
        machine.openshift.io/cluster-api-machine-type: worker
        machine.openshift.io/cluster-api-machineset: ostest-hwmdt-worker-0
    spec:
      metadata: {}
      providerSpec:
        value:
          apiVersion: baremetal.cluster.k8s.io/v1alpha1
          hostSelector: {}
          image:
            checksum: http://172.22.0.3:6181/images/rhcos-<version>.<architecture>.qcow2.
            <md5sum> 1
            url: http://172.22.0.3:6181/images/rhcos-<version>.<architecture>.qcow2 2
          kind: BareMetalMachineProviderSpec
          metadata:
            creationTimestamp: null
          userData:
            name: worker-user-data
      status:
        availableReplicas: 2
        fullyLabeledReplicas: 2
        observedGeneration: 11
        readyReplicas: 2
        replicas: 2
```

- Edit the **checksum** URL to use the API VIP address.

- 2 Edit the **url** URL to use the API VIP address.

16.5.4. Diagnosing a duplicate MAC address when provisioning a new host in the cluster

If the MAC address of an existing bare-metal node in the cluster matches the MAC address of a bare-metal host you are attempting to add to the cluster, the Bare Metal Operator associates the host with the existing node. If the host enrollment, inspection, cleaning, or other Ironic steps fail, the Bare Metal Operator retries the installation continuously. A registration error is displayed for the failed bare-metal host.

You can diagnose a duplicate MAC address by examining the bare-metal hosts that are running in the **openshift-machine-api** namespace.

Prerequisites

- Install an OpenShift Container Platform cluster on bare metal.
- Install the OpenShift Container Platform CLI **oc**.
- Log in as a user with **cluster-admin** privileges.

Procedure

To determine whether a bare-metal host that fails provisioning has the same MAC address as an existing node, do the following:

1. Get the bare-metal hosts running in the **openshift-machine-api** namespace:

```
$ oc get bmh -n openshift-machine-api
```

Example output

NAME	STATUS	PROVISIONING STATUS	CONSUMER
openshift-master-0	OK	externally provisioned	openshift-zpwpq-master-0
openshift-master-1	OK	externally provisioned	openshift-zpwpq-master-1
openshift-master-2	OK	externally provisioned	openshift-zpwpq-master-2
openshift-worker-0	OK	provisioned	openshift-zpwpq-worker-0-lv84n
openshift-worker-1	OK	provisioned	openshift-zpwpq-worker-0-zd8lm
openshift-worker-2	error	registering	

2. To see more detailed information about the status of the failing host, run the following command replacing **<bare_metal_host_name>** with the name of the host:

```
$ oc get -n openshift-machine-api bmh <bare_metal_host_name> -o yaml
```

Example output

```
...
status:
  errorCount: 12
  errorMessage: MAC address b4:96:91:1d:7c:20 conflicts with existing node openshift-
```

```
worker-1
  errorType: registration error
  ...
```

16.5.5. Provisioning the bare metal node

Provisioning the bare metal node requires executing the following procedure from the provisioner node.

Procedure

1. Ensure the **STATE** is **available** before provisioning the bare metal node.

```
$ oc -n openshift-machine-api get bmh openshift-worker-<num>
```

Where **<num>** is the worker node number.

```
NAME           STATE    ONLINE ERROR AGE
openshift-worker  available true     34h
```

2. Get a count of the number of worker nodes.

```
$ oc get nodes
```

```
NAME                                                    STATUS  ROLES    AGE  VERSION
openshift-master-1.openshift.example.com              Ready   master   30h  v1.29.4
openshift-master-2.openshift.example.com              Ready   master   30h  v1.29.4
openshift-master-3.openshift.example.com              Ready   master   30h  v1.29.4
openshift-worker-0.openshift.example.com              Ready   worker   30h  v1.29.4
openshift-worker-1.openshift.example.com              Ready   worker   30h  v1.29.4
```

3. Get the compute machine set.

```
$ oc get machinesets -n openshift-machine-api
```

```
NAME           DESIRED  CURRENT  READY  AVAILABLE  AGE
...
openshift-worker-0.example.com  1        1        1        1          55m
openshift-worker-1.example.com  1        1        1        1          55m
```

4. Increase the number of worker nodes by one.

```
$ oc scale --replicas=<num> machineset <machineset> -n openshift-machine-api
```

Replace **<num>** with the new number of worker nodes. Replace **<machineset>** with the name of the compute machine set from the previous step.

5. Check the status of the bare metal node.

```
$ oc -n openshift-machine-api get bmh openshift-worker-<num>
```

Where **<num>** is the worker node number. The STATE changes from **ready** to **provisioning**.

NAME	STATE	CONSUMER	ONLINE	ERROR
openshift-worker-<num>	provisioning	openshift-worker-<num>-65tjz		true

The **provisioning** status remains until the OpenShift Container Platform cluster provisions the node. This can take 30 minutes or more. After the node is provisioned, the state will change to **provisioned**.

NAME	STATE	CONSUMER	ONLINE	ERROR
openshift-worker-<num>	provisioned	openshift-worker-<num>-65tjz		true

- After provisioning completes, ensure the bare metal node is ready.

```
$ oc get nodes
```

NAME	STATUS	ROLES	AGE	VERSION
openshift-master-1.openshift.example.com	Ready	master	30h	v1.29.4
openshift-master-2.openshift.example.com	Ready	master	30h	v1.29.4
openshift-master-3.openshift.example.com	Ready	master	30h	v1.29.4
openshift-worker-0.openshift.example.com	Ready	worker	30h	v1.29.4
openshift-worker-1.openshift.example.com	Ready	worker	30h	v1.29.4
openshift-worker-<num>.openshift.example.com	Ready	worker	3m27s	v1.29.4

You can also check the kubelet.

```
$ ssh openshift-worker-<num>
```

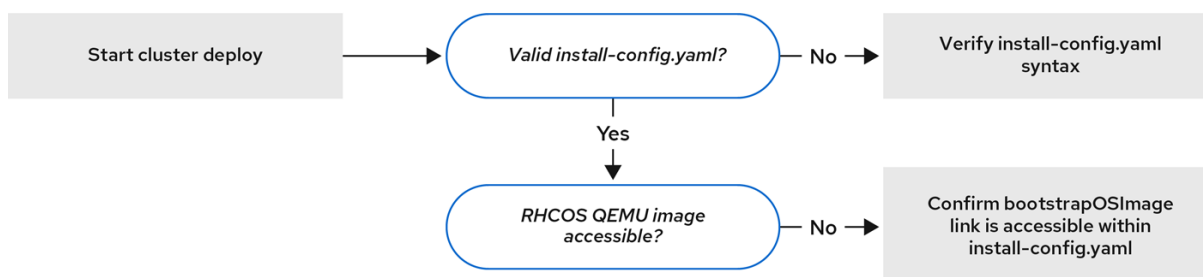
```
[kni@openshift-worker-<num>]$ journalctl -fu kubelet
```

16.6. TROUBLESHOOTING

16.6.1. Troubleshooting the installation program workflow

Prior to troubleshooting the installation environment, it is critical to understand the overall flow of the installer-provisioned installation on bare metal. The diagrams below provide a troubleshooting flow with a step-by-step breakdown for the environment.

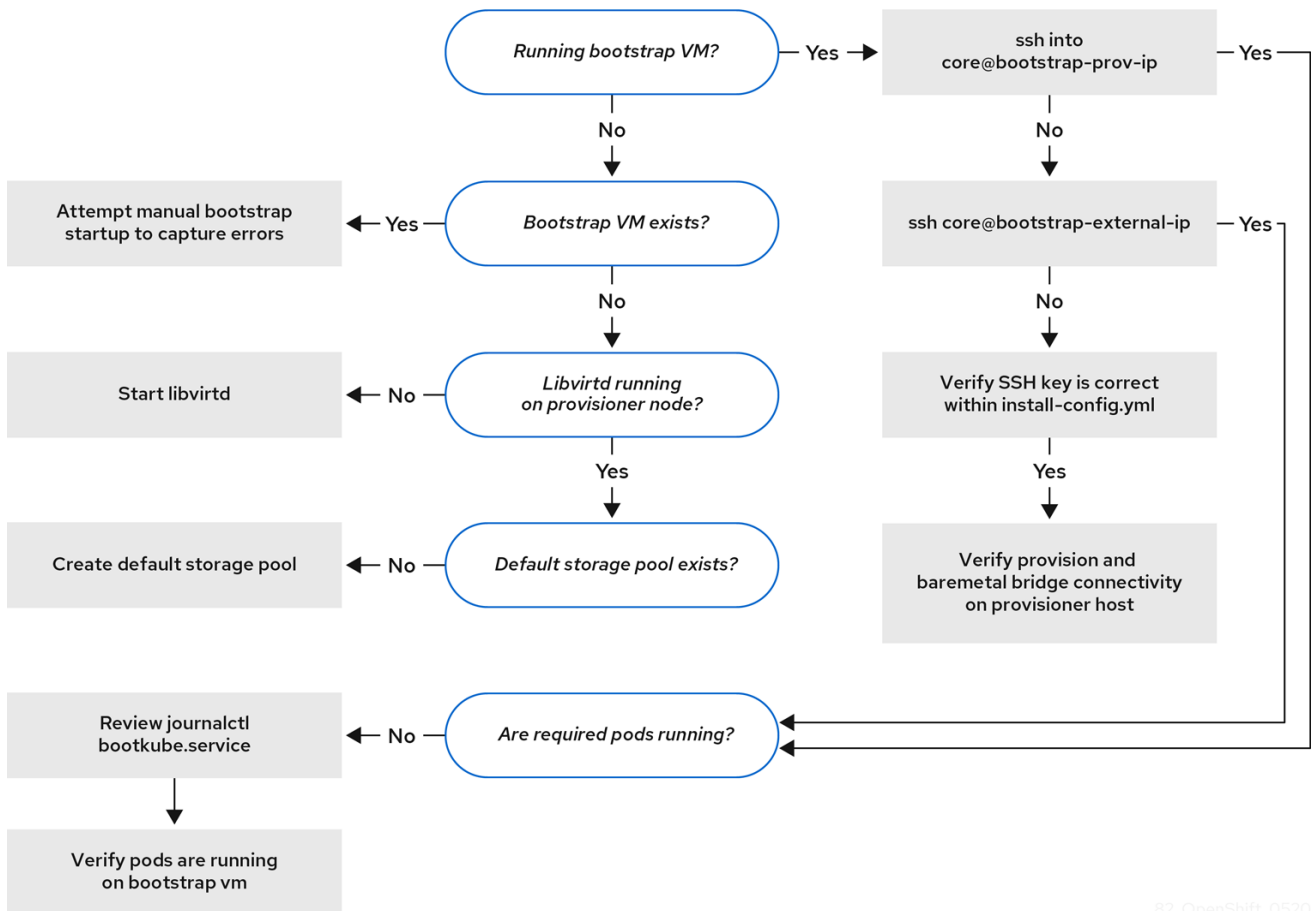
Workflow 1 of 4



82_OpenShift_0420

Workflow 1 of 4 illustrates a troubleshooting workflow when the **install-config.yaml** file has errors or the Red Hat Enterprise Linux CoreOS (RHCOS) images are inaccessible. Troubleshooting suggestions can be found at [Troubleshooting install-config.yaml](#).

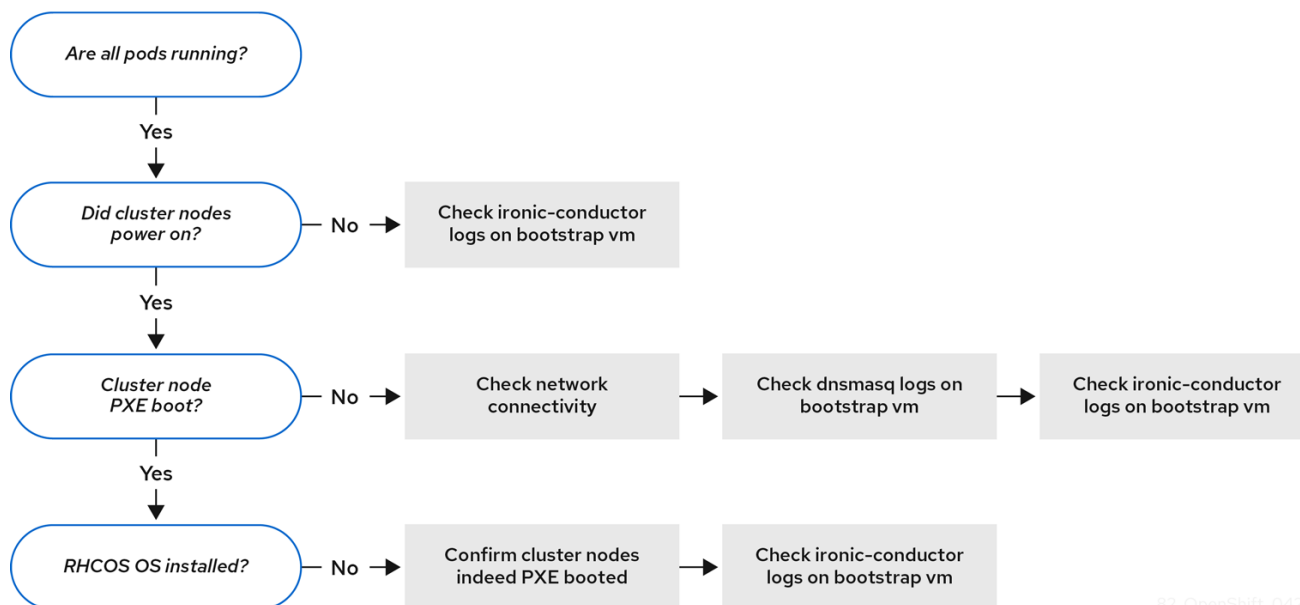
Workflow 2 of 4



82_OpenShift_0520

Workflow 2 of 4 illustrates a troubleshooting workflow for [bootstrap VM issues](#), [bootstrap VMs that cannot boot up the cluster nodes](#), and [inspecting logs](#). When installing an OpenShift Container Platform cluster without the **provisioning** network, this workflow does not apply.

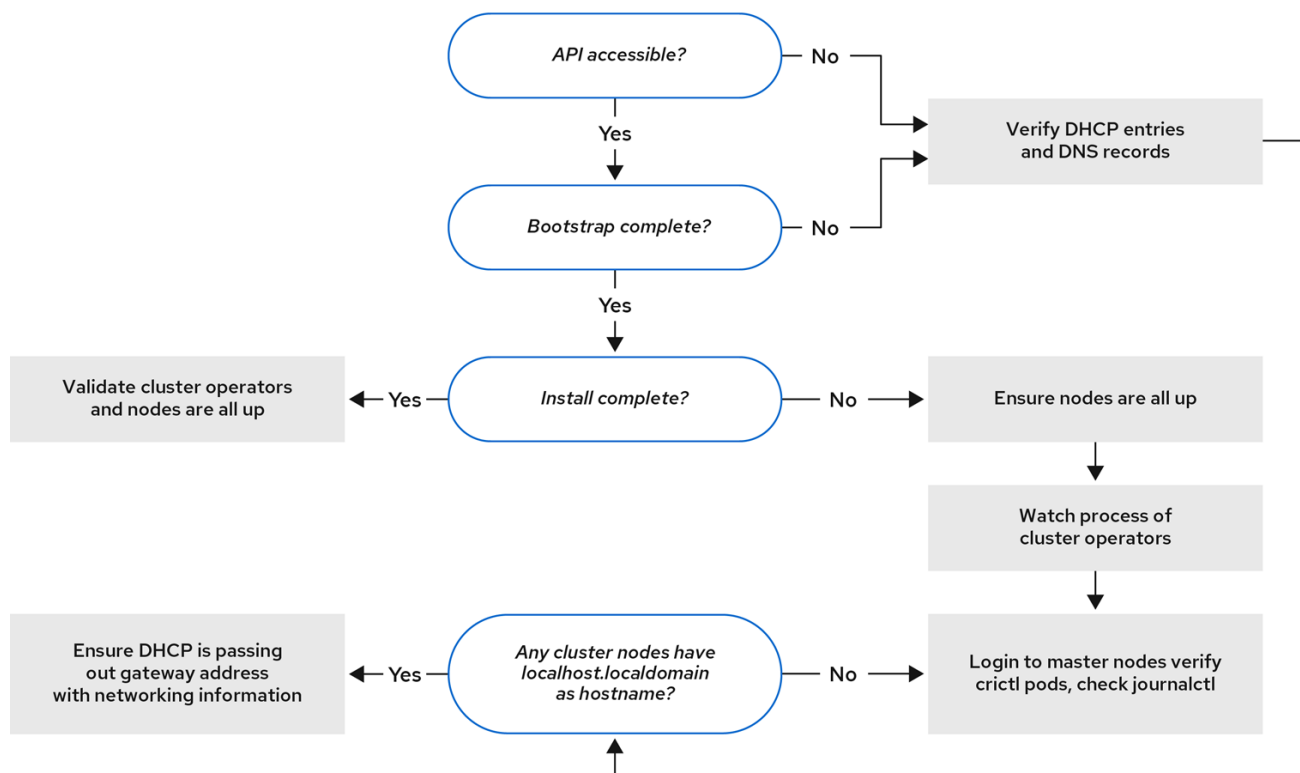
Workflow 3 of 4



82_OpenShift_0420

Workflow 3 of 4 illustrates a troubleshooting workflow for [cluster nodes that will not PXE boot](#). If installing using RedFish Virtual Media, each node must meet minimum firmware requirements for the installation program to deploy the node. See [Firmware requirements for installing with virtual media](#) in the [Prerequisites](#) section for additional details.

Workflow 4 of 4



82_OpenShift_0420

Workflow 4 of 4 illustrates a troubleshooting workflow from [a non-accessible API](#) to a [validated installation](#).

16.6.2. Troubleshooting install-config.yaml

The **install-config.yaml** configuration file represents all of the nodes that are part of the OpenShift Container Platform cluster. The file contains the necessary options consisting of but not limited to **apiVersion**, **baseDomain**, **imageContentSources** and virtual IP addresses. If errors occur early in the deployment of the OpenShift Container Platform cluster, the errors are likely in the **install-config.yaml** configuration file.

Procedure

1. Use the guidelines in [YAML-tips](#).
2. Verify the YAML syntax is correct using [syntax-check](#).
3. Verify the Red Hat Enterprise Linux CoreOS (RHCOS) QEMU images are properly defined and accessible via the URL provided in the **install-config.yaml**. For example:

```
$ curl -s -o /dev/null -I -w "%{http_code}\n" http://webserver.example.com:8080/rhcos-44.81.202004250133-0-qemu.<architecture>.qcow2.gz?sha256=7d884b46ee54fe87bbc3893bf2aa99af3b2d31f2e19ab5529c60636fbd0f1ce7
```

If the output is **200**, there is a valid response from the webserver storing the bootstrap VM image.

16.6.3. Troubleshooting bootstrap VM issues


The OpenShift Container Platform installation program spawns a bootstrap node virtual machine, which handles provisioning the OpenShift Container Platform cluster nodes.

Procedure

1. About 10 to 15 minutes after triggering the installation program, check to ensure the bootstrap VM is operational using the **virsh** command:

```
$ sudo virsh list
```

```
Id Name State
-----
12 openshift-xf6fq-bootstrap running
```



NOTE

The name of the bootstrap VM is always the cluster name followed by a random set of characters and ending in the word "bootstrap."

2. If the bootstrap VM is not running after 10-15 minutes, verify **libvirt** is running on the system by executing the following command:

```
$ systemctl status libvirt
```

- `libvirtd.service` - Virtualization daemon
 Loaded: loaded (/usr/lib/systemd/system/libvirtd.service; enabled; vendor preset: enabled)
 Active: active (running) since Tue 2020-03-03 21:21:07 UTC; 3 weeks 5 days ago
 Docs: man:libvirtd(8)
 <https://libvirt.org>
 Main PID: 9850 (libvirtd)
 Tasks: 20 (limit: 32768)
 Memory: 74.8M
 CGroup: /system.slice/libvirtd.service
 └─ 9850 /usr/sbin/libvirtd

If the bootstrap VM is operational, log in to it.

3. Use the **virsh console** command to find the IP address of the bootstrap VM:

```
$ sudo virsh console example.com
```

```
Connected to domain example.com
Escape character is ^]
Red Hat Enterprise Linux CoreOS 43.81.202001142154.0 (Ootpa) 4.3
SSH host key: SHA256:BRWJktXZgQQRY5zjuAV0IKZ4WM7i4TiUyMVanqu9Pqg (ED25519)
SSH host key: SHA256:7+iKGA7VtG5szmk2jB5gl/5EZ+SNcJ3a2g23o0Inlio (ECDSA)
SSH host key: SHA256:DH5VWhvhvagOTaLsYiVNse9ca+ZSW/30OOMed8rlGOc (RSA)
ens3: fd35:919d:4042:2:c7ed:9a9f:a9ec:7
ens4: 172.22.0.2 fe80::1d05:e52e:be5d:263f
localhost login:
```



IMPORTANT

When deploying an OpenShift Container Platform cluster without the **provisioning** network, you must use a public IP address and not a private IP address like **172.22.0.2**.

4. After you obtain the IP address, log in to the bootstrap VM using the **ssh** command:



NOTE

In the console output of the previous step, you can use the IPv6 IP address provided by **ens3** or the IPv4 IP provided by **ens4**.

```
$ ssh core@172.22.0.2
```

If you are not successful logging in to the bootstrap VM, you have likely encountered one of the following scenarios:

- You cannot reach the **172.22.0.0/24** network. Verify the network connectivity between the provisioner and the **provisioning** network bridge. This issue might occur if you are using a **provisioning** network.
- You cannot reach the bootstrap VM through the public network. When attempting to SSH via **baremetal** network, verify connectivity on the **provisioner** host specifically around the **baremetal** network bridge.

- You encountered **Permission denied (publickey,password,keyboard-interactive)**. When attempting to access the bootstrap VM, a **Permission denied** error might occur. Verify that the SSH key for the user attempting to log in to the VM is set within the **install-config.yaml** file.

16.6.3.1. Bootstrap VM cannot boot up the cluster nodes

During the deployment, it is possible for the bootstrap VM to fail to boot the cluster nodes, which prevents the VM from provisioning the nodes with the RHCOS image. This scenario can arise due to:

- A problem with the **install-config.yaml** file.
- Issues with out-of-band network access when using the baremetal network.

To verify the issue, there are three containers related to **ironic**:

- **ironic**
- **ironic-inspector**

Procedure

1. Log in to the bootstrap VM:

```
$ ssh core@172.22.0.2
```

2. To check the container logs, execute the following:

```
[core@localhost ~]$ sudo podman logs -f <container_name>
```

Replace **<container_name>** with one of **ironic** or **ironic-inspector**. If you encounter an issue where the control plane nodes are not booting up from PXE, check the **ironic** pod. The **ironic** pod contains information about the attempt to boot the cluster nodes, because it attempts to log in to the node over IPMI.

Potential reason

The cluster nodes might be in the **ON** state when deployment started.

Solution

Power off the OpenShift Container Platform cluster nodes before you begin the installation over IPMI:

```
$ ipmitool -I lanplus -U root -P <password> -H <out_of_band_ip> power off
```

16.6.3.2. Inspecting logs

When experiencing issues downloading or accessing the RHCOS images, first verify that the URL is correct in the **install-config.yaml** configuration file.

Example of internal webserver hosting RHCOS images

```
bootstrapOSImage: http://<ip:port>/rhcos-43.81.202001142154.0-qemu.<architecture>.qcow2.gz?  
sha256=9d999f55ff1d44f7ed7c106508e5deecd04dc3c06095d34d36bf1cd127837e0c  
clusterOSImage: http://<ip:port>/rhcos-43.81.202001142154.0-openstack.<architecture>.qcow2.gz?  
sha256=a1bda656fa0892f7b936fdc6b6a6086bddaed5dafacedcd7a1e811abb78fe3b0
```

The **coreos-downloader** container downloads resources from a webserver or from the external quay.io registry, whichever the **install-config.yaml** configuration file specifies. Verify that the **coreos-downloader** container is up and running and inspect its logs as needed.

Procedure

1. Log in to the bootstrap VM:

```
$ ssh core@172.22.0.2
```

2. Check the status of the **coreos-downloader** container within the bootstrap VM by running the following command:

```
[core@localhost ~]$ sudo podman logs -f coreos-downloader
```

If the bootstrap VM cannot access the URL to the images, use the **curl** command to verify that the VM can access the images.

3. To inspect the **bootkube** logs that indicate if all the containers launched during the deployment phase, execute the following:

```
[core@localhost ~]$ journalctl -xe
```

```
[core@localhost ~]$ journalctl -b -f -u bootkube.service
```

4. Verify all the pods, including **dnsmasq**, **mariadb**, **httpd**, and **ironic**, are running:

```
[core@localhost ~]$ sudo podman ps
```

5. If there are issues with the pods, check the logs of the containers with issues. To check the logs of the **ironic** service, run the following command:

```
[core@localhost ~]$ sudo podman logs ironic
```

16.6.4. Investigating an unavailable Kubernetes API

When the Kubernetes API is unavailable, check the control plane nodes to ensure that they are running the correct components. Also, check the hostname resolution.

Procedure

1. Ensure that **etcd** is running on each of the control plane nodes by running the following command:

```
$ sudo crictl logs $(sudo crictl ps --pod=$(sudo crictl pods --name=etcd-member --quiet) --quiet)
```

2. If the previous command fails, ensure that Kubelet created the **etcd** pods by running the following command:

```
$ sudo crictl pods --name=etcd-member
```

If there are no pods, investigate **etcd**.

3. Check the cluster nodes to ensure they have a fully qualified domain name, and not just **localhost.localdomain**, by using the following command:

```
$ hostname
```

If a hostname is not set, set the correct hostname. For example:

```
$ sudo hostnamectl set-hostname <hostname>
```

4. Ensure that each node has the correct name resolution in the DNS server using the **dig** command:

```
$ dig api.<cluster_name>.example.com
```

```
;; <<>> DiG 9.11.4-P2-RedHat-9.11.4-26.P2.el8 <<>> api.<cluster_name>.example.com
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 37551
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 1, ADDITIONAL: 2

;; OPT PSEUDOSECTION:
;; EDNS: version: 0, flags:; udp: 4096
;; COOKIE: 866929d2f8e8563582af23f05ec44203d313e50948d43f60 (good)
;; QUESTION SECTION:
;api.<cluster_name>.example.com. IN A

;; ANSWER SECTION:
api.<cluster_name>.example.com. 10800 IN A 10.19.13.86

;; AUTHORITY SECTION:
<cluster_name>.example.com. 10800 IN NS <cluster_name>.example.com.

;; ADDITIONAL SECTION:
<cluster_name>.example.com. 10800 IN A 10.19.14.247

;; Query time: 0 msec
;; SERVER: 10.19.14.247#53(10.19.14.247)
;; WHEN: Tue May 19 20:30:59 UTC 2020
;; MSG SIZE rcvd: 140
```

The output in the foregoing example indicates that the appropriate IP address for the **api.<cluster_name>.example.com** VIP is **10.19.13.86**. This IP address should reside on the **baremetal** network.

16.6.5. Troubleshooting a failure to initialize the cluster

The installation program uses the Cluster Version Operator to create all the components of an OpenShift Container Platform cluster. When the installation program fails to initialize the cluster, you can retrieve the most important information from the **ClusterVersion** and **ClusterOperator** objects.

Procedure

1. Inspect the **ClusterVersion** object by running the following command:

```
$ oc --kubeconfig=${INSTALL_DIR}/auth/kubeconfig get clusterversion -o yaml
```

Example output

```
apiVersion: config.openshift.io/v1
kind: ClusterVersion
metadata:
  creationTimestamp: 2019-02-27T22:24:21Z
  generation: 1
  name: version
  resourceVersion: "19927"
  selfLink: /apis/config.openshift.io/v1/clusterversions/version
  uid: 6e0f4cf8-3ade-11e9-9034-0a923b47ded4
spec:
  channel: stable-4.1
  clusterID: 5ec312f9-f729-429d-a454-61d4906896ca
status:
  availableUpdates: null
  conditions:
  - lastTransitionTime: 2019-02-27T22:50:30Z
    message: Done applying 4.1.1
    status: "True"
    type: Available
  - lastTransitionTime: 2019-02-27T22:50:30Z
    status: "False"
    type: Failing
  - lastTransitionTime: 2019-02-27T22:50:30Z
    message: Cluster version is 4.1.1
    status: "False"
    type: Progressing
  - lastTransitionTime: 2019-02-27T22:24:31Z
    message: 'Unable to retrieve available updates: unknown version 4.1.1
    reason: RemoteFailed
    status: "False"
    type: RetrievedUpdates
  desired:
    image: registry.svc.ci.openshift.org/openshift/origin-
release@sha256:91e6f754975963e7db1a9958075eb609ad226968623939d262d1cf45e9dbc3
9a
    version: 4.1.1
  history:
  - completionTime: 2019-02-27T22:50:30Z
    image: registry.svc.ci.openshift.org/openshift/origin-
release@sha256:91e6f754975963e7db1a9958075eb609ad226968623939d262d1cf45e9dbc3
9a
    startedTime: 2019-02-27T22:24:31Z
    state: Completed
    version: 4.1.1
  observedGeneration: 1
  versionHash: Wa7as_ik1qE=
```

2. View the conditions by running the following command:

```
$ oc --kubeconfig=${INSTALL_DIR}/auth/kubeconfig get clusterversion version \
-o=jsonpath='{range .status.conditions[*]}{.type}" "{.status}" "{.message}"{"\n"}{end}'
```

Some of most important conditions include **Failing**, **Available** and **Progressing**.

Example output

```
Available True Done applying 4.1.1
Failing False
Progressing False Cluster version is 4.0.0-0.alpha-2019-02-26-194020
RetrievedUpdates False Unable to retrieve available updates: unknown version 4.1.1
```

- Inspect the **ClusterOperator** object by running the following command:

```
$ oc --kubeconfig=${INSTALL_DIR}/auth/kubeconfig get clusteroperator
```

The command returns the status of the cluster Operators.

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	FAILING	SINCE
cluster-baremetal-operator		True	False	False	17m
cluster-autoscaler		True	False	False	17m
cluster-storage-operator		True	False	False	10m
console	True	False	False		7m21s
dns	True	False	False		31m
image-registry	True	False	False		9m58s
ingress	True	False	False		10m
kube-apiserver	True	False	False		28m
kube-controller-manager	True	False	False		21m
kube-scheduler	True	False	False		25m
machine-api	True	False	False		17m
machine-config	True	False	False		17m
marketplace-operator	True	False	False		10m
monitoring	True	False	False		8m23s
network	True	False	False		13m
node-tuning	True	False	False		11m
openshift-apiserver	True	False	False		15m
openshift-authentication	True	False	False		20m
openshift-cloud-credential-operator	True	False	False		18m
openshift-controller-manager	True	False	False		10m
openshift-samples	True	False	False		8m42s
operator-lifecycle-manager	True	False	False		17m
service-ca	True	False	False		30m

- Inspect individual cluster Operators by running the following command:

```
$ oc --kubeconfig=${INSTALL_DIR}/auth/kubeconfig get clusteroperator <operator> -oyaml
```

1

- 1 Replace **<operator>** with the name of a cluster Operator. This command is useful for identifying why an cluster Operator has not achieved the **Available** state or is in the **Failed** state.

Example output

```

apiVersion: config.openshift.io/v1
kind: ClusterOperator
metadata:
  creationTimestamp: 2019-02-27T22:47:04Z
  generation: 1
  name: monitoring
  resourceVersion: "24677"
  selfLink: /apis/config.openshift.io/v1/clusteroperators/monitoring
  uid: 9a6a5ef9-3ae1-11e9-bad4-0a97b6ba9358
spec: {}
status:
  conditions:
  - lastTransitionTime: 2019-02-27T22:49:10Z
    message: Successfully rolled out the stack.
    status: "True"
    type: Available
  - lastTransitionTime: 2019-02-27T22:49:10Z
    status: "False"
    type: Progressing
  - lastTransitionTime: 2019-02-27T22:49:10Z
    status: "False"
    type: Failing
  extension: null
  relatedObjects: null
  version: ""

```

- To get the cluster Operator's status condition, run the following command:

```

$ oc --kubeconfig=${INSTALL_DIR}/auth/kubeconfig get clusteroperator <operator> \
  -o=jsonpath='{range .status.conditions[*]}{.type}{" "}{.status}{" "}{.message}{"\n"}{end}'

```

Replace **<operator>** with the name of one of the operators above.

Example output

```

Available True Successfully rolled out the stack
Progressing False
Failing False

```

- To retrieve the list of objects owned by the cluster Operator, execute the following command:

```

oc --kubeconfig=${INSTALL_DIR}/auth/kubeconfig get clusteroperator kube-apiserver \
  -o=jsonpath='{.status.relatedObjects}'

```

Example output

```

[map[resource:kubeapiservers group:operator.openshift.io name:cluster] map[group:
name:openshift-config resource:namespaces] map[group: name:openshift-config-managed
resource:namespaces] map[group: name:openshift-kube-apiserver-operator
resource:namespaces] map[group: name:openshift-kube-apiserver resource:namespaces]]

```

16.6.6. Troubleshooting a failure to fetch the console URL

The installation program retrieves the URL for the OpenShift Container Platform console by using `[route][route-object]` within the `openshift-console` namespace. If the installation program fails to retrieve the URL for the console, use the following procedure.

Procedure

1. Check if the console router is in the **Available** or **Failing** state by running the following command:

```
$ oc --kubeconfig=${INSTALL_DIR}/auth/kubeconfig get clusteroperator console -oyaml
```

```
apiVersion: config.openshift.io/v1
kind: ClusterOperator
metadata:
  creationTimestamp: 2019-02-27T22:46:57Z
  generation: 1
  name: console
  resourceVersion: "19682"
  selfLink: /apis/config.openshift.io/v1/clusteroperators/console
  uid: 960364aa-3ae1-11e9-bad4-0a97b6ba9358
spec: {}
status:
  conditions:
  - lastTransitionTime: 2019-02-27T22:46:58Z
    status: "False"
    type: Failing
  - lastTransitionTime: 2019-02-27T22:50:12Z
    status: "False"
    type: Progressing
  - lastTransitionTime: 2019-02-27T22:50:12Z
    status: "True"
    type: Available
  - lastTransitionTime: 2019-02-27T22:46:57Z
    status: "True"
    type: Upgradeable
  extension: null
  relatedObjects:
  - group: operator.openshift.io
    name: cluster
    resource: consoles
  - group: config.openshift.io
    name: cluster
    resource: consoles
  - group: oauth.openshift.io
    name: console
    resource: oauthclients
  - group: ""
    name: openshift-console-operator
    resource: namespaces
  - group: ""
    name: openshift-console
    resource: namespaces
  versions: null
```

2. Manually retrieve the console URL by executing the following command:

```
$ oc --kubeconfig=${INSTALL_DIR}/auth/kubeconfig get route console -n openshift-console \
  -o=jsonpath='{.spec.host}' console-openshift-console.apps.adahiya-
  1.devcluster.openshift.com
```

16.6.7. Troubleshooting a failure to add the ingress certificate to kubeconfig

The installation program adds the default ingress certificate to the list of trusted client certificate authorities in `${INSTALL_DIR}/auth/kubeconfig`. If the installation program fails to add the ingress certificate to the `kubeconfig` file, you can retrieve the certificate from the cluster and add it.

Procedure

1. Retrieve the certificate from the cluster using the following command:

```
$ oc --kubeconfig=${INSTALL_DIR}/auth/kubeconfig get configmaps default-ingress-cert \
  -n openshift-config-managed -o=jsonpath='{.data.ca-bundle.crt}'
```

```
-----BEGIN CERTIFICATE-----
MIIC/TCCAeWgAwIBAgIBATANBgkqhkiG9w0BAQsFADAuMSwwKgYDVQQDDCNjbHVz
dGVyLWluZ3Jlc3MtYm93LWVudC00b3JAMTU1MTMwNzU0OTAEFw0xOTAyMjcyMjMjha
Fw0yMTAyMjYyMjMjMjMjMC4xLDAqBgNVBAMMI2NsdXN0ZXItaW5ncmVzcy1vcGVy
YXRvckAxNTUxMzA3NTg5MIIIBjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEA
uCA4fQ+2YXoXSUL4h/mcvJfrgpBfKBW5hfB8NcgXeCYiQPnCKbIH1sEQnl3VC5Pk
2OfNCF3PUlfm4i8CHC95a7nChRjMjNg1gVrWCvS/ohLgnO0BvszSiRLxlpuo3C4S
EVqqvxValHcbdAXWgZLQoYZXV7RMz8yZjI5CfhDaaltyBFj3GtIJKXgUwp/5sUfl
LDXW8MM6AXfuG+kweLdLCMm3g8WLLfLbLvVBKB+4IhIH7II0buOz04RKhnYN+Ebw
tcvFi55vwuUCWMnGhWHGEQ8sWm/wLnNIOwsUz7S1/sW8nj87GFHzgkaVM9EOnoNI
gKhMBK9ItNzjrP6dgiKBCQIDAQABoyYwJDAOBgNVHQ8BAf8EBAMCAQwEgYDVR0T
AQH/BAgwbGEB/wIBADANBgkqhkiG9w0BAQsFAAOCAQEAg+vi0sFKudaZ9aUQMMha
CeWx9CZvZBblnAWT/61UdpZKpFi4eJ2d33IGcfKwHOi2NP/iSKQBebfG0iNLVVPz
vwLbSG1i9R9GLdAbnHpPT9UG6fLaDloKpnKiBfGENfxeiq5vTln2bAgivxrVlyiq
+MdDXFAwb6V4u2xh6RChI7akNsS3oU9PZ9YOs5e8vJp2YAEphht05X0swA+X8V8T
C278FFifpo0h3Q0Dbv8Rfn4UpBEtN4KkLeS+JeT+0o2XOsFZp7Uhr9yFlodRsnNo
H/Uwmab28ocNrGNiEVaVH6eTTQeeZuOdoQzUbCIElpVmkrNGY0M42K0PvOQ/e7+y
AQ==
-----END CERTIFICATE-----
```

2. Add the certificate to the `client-certificate-authority-data` field in the `${INSTALL_DIR}/auth/kubeconfig` file.

16.6.8. Troubleshooting SSH access to cluster nodes

For added security, you cannot SSH into the cluster from outside the cluster by default. However, you can access control plane and worker nodes from the provisioner node. If you cannot SSH into the cluster nodes from the provisioner node, the nodes might be waiting on the bootstrap VM. The control plane nodes retrieve their boot configuration from the bootstrap VM, and they cannot boot successfully if they do not retrieve the boot configuration.

Procedure

1. If you have physical access to the nodes, check their console output to determine if they have successfully booted. If the nodes are still retrieving their boot configuration, there might be problems with the bootstrap VM .
2. Ensure you have configured the **sshKey: '<ssh_pub_key>'** setting in the **install-config.yaml** file, where **<ssh_pub_key>** is the public key of the **kni** user on the provisioner node.

16.6.9. Cluster nodes will not PXE boot

When OpenShift Container Platform cluster nodes will not PXE boot, execute the following checks on the cluster nodes that will not PXE boot. This procedure does not apply when installing an OpenShift Container Platform cluster without the **provisioning** network.

Procedure

1. Check the network connectivity to the **provisioning** network.
2. Ensure PXE is enabled on the NIC for the **provisioning** network and PXE is disabled for all other NICs.
3. Verify that the **install-config.yaml** configuration file includes the **rootDeviceHints** parameter and boot MAC address for the NIC connected to the **provisioning** network. For example:

control plane node settings

```
bootMACAddress: 24:6E:96:1B:96:90 # MAC of bootable provisioning NIC
```

Worker node settings

```
bootMACAddress: 24:6E:96:1B:96:90 # MAC of bootable provisioning NIC
```

16.6.10. Installing creates no worker nodes

The installation program does not provision worker nodes directly. Instead, the Machine API Operator scales nodes up and down on supported platforms. If worker nodes are not created after 15 to 20 minutes, depending on the speed of the cluster's internet connection, investigate the Machine API Operator.

Procedure

1. Check the Machine API Operator by running the following command:

```
$ oc --kubeconfig=${INSTALL_DIR}/auth/kubeconfig \
  --namespace=openshift-machine-api get deployments
```

If **\${INSTALL_DIR}** is not set in your environment, replace the value with the name of the installation directory.

Example output

NAME	READY	UP-TO-DATE	AVAILABLE	AGE
cluster-autoscaler-operator	1/1	1	1	86m
cluster-baremetal-operator	1/1	1	1	86m

```

machine-api-controllers 1/1 1 1 85m
machine-api-operator 1/1 1 1 86m

```

2. Check the machine controller logs by running the following command:

```

$ oc --kubeconfig=${INSTALL_DIR}/auth/kubeconfig \
  --namespace=openshift-machine-api logs deployments/machine-api-controllers \
  --container=machine-controller

```

16.6.11. Troubleshooting the Cluster Network Operator

The Cluster Network Operator is responsible for deploying the networking components. It runs early in the installation process, after the control plane nodes have come up but before the installation program removes the bootstrap control plane. Issues with this Operator might indicate installation program issues.

Procedure

1. Ensure the network configuration exists by running the following command:

```
$ oc get network -o yaml cluster
```

If it does not exist, the installation program did not create it. To find out why, run the following command:

```
$ openshift-install create manifests
```

Review the manifests to determine why the installation program did not create the network configuration.

2. Ensure the network is running by entering the following command:

```
$ oc get po -n openshift-network-operator
```

16.6.12. Unable to discover new bare metal hosts using the BMC

In some cases, the installation program will not be able to discover the new bare metal hosts and issue an error, because it cannot mount the remote virtual media share.

For example:

```

ProvisioningError 51s metal3-baremetal-controller Image provisioning failed: Deploy step
deploy.deploy failed with BadRequestError: HTTP POST
https://<bmc_address>/redfish/v1/Managers/iDRAC.Embedded.1/VirtualMedia/CD/Actions/VirtualMedia.
InsertMedia
returned code 400.
Base.1.8.GeneralError: A general error has occurred. See ExtendedInfo for more information
Extended information: [
{
  "Message": "Unable to mount remote share https://<ironic_address>/redfish/boot-<uuid>.iso.",
  "MessageArgs": [
    "https://<ironic_address>/redfish/boot-<uuid>.iso"
  ],
}
]

```

```

    "MessageArgs@odata.count": 1,
    "MessageId": "IDRAC.2.5.RAC0720",
    "RelatedProperties": [
      "#/Image"
    ],
    "RelatedProperties@odata.count": 1,
    "Resolution": "Retry the operation.",
    "Severity": "Informational"
  }
].

```

In this situation, if you are using virtual media with an unknown certificate authority, you can configure your baseboard management controller (BMC) remote file share settings to trust an unknown certificate authority to avoid this error.



NOTE

This resolution was tested on OpenShift Container Platform 4.11 with Dell iDRAC 9 and firmware version 5.10.50.

16.6.13. Troubleshooting worker nodes that cannot join the cluster

Installer-provisioned clusters deploy with a DNS server that includes a DNS entry for the **api-int.<cluster_name>.<base_domain>** URL. If the nodes within the cluster use an external or upstream DNS server to resolve the **api-int.<cluster_name>.<base_domain>** URL and there is no such entry, worker nodes might fail to join the cluster. Ensure that all nodes in the cluster can resolve the domain name.

Procedure

1. Add a DNS A/AAAA or CNAME record to internally identify the API load balancer. For example, when using dnsmasq, modify the **dnsmasq.conf** configuration file:

```
$ sudo nano /etc/dnsmasq.conf
```

```

address=/api-int.<cluster_name>.<base_domain>/<IP_address>
address=/api-int.mycluster.example.com/192.168.1.10
address=/api-int.mycluster.example.com/2001:0db8:85a3:0000:0000:8a2e:0370:7334

```

2. Add a DNS PTR record to internally identify the API load balancer. For example, when using dnsmasq, modify the **dnsmasq.conf** configuration file:

```
$ sudo nano /etc/dnsmasq.conf
```

```

ptr-record=<IP_address>.in-addr.arpa,api-int.<cluster_name>.<base_domain>
ptr-record=10.1.168.192.in-addr.arpa,api-int.mycluster.example.com

```

3. Restart the DNS server. For example, when using dnsmasq, execute the following command:

```
$ sudo systemctl restart dnsmasq
```

These records must be resolvable from all the nodes within the cluster.

16.6.14. Cleaning up previous installations

In the event of a previous failed deployment, remove the artifacts from the failed attempt before attempting to deploy OpenShift Container Platform again.

Procedure

1. Power off all bare metal nodes prior to installing the OpenShift Container Platform cluster:

```
$ ipmitool -I lanplus -U <user> -P <password> -H <management_server_ip> power off
```

2. Remove all old bootstrap resources if any are left over from a previous deployment attempt:

```
for i in $(sudo virsh list | tail -n +3 | grep bootstrap | awk {'print $2'});
do
  sudo virsh destroy $i;
  sudo virsh undefine $i;
  sudo virsh vol-delete $i --pool $i;
  sudo virsh vol-delete $i.ign --pool $i;
  sudo virsh pool-destroy $i;
  sudo virsh pool-undefine $i;
done
```


16.6.15. Issues with creating the registry

When creating a disconnected registry, you might encounter a "User Not Authorized" error when attempting to mirror the registry. This error might occur if you fail to append the new authentication to the existing **pull-secret.txt** file.

Procedure

1. Check to ensure authentication is successful:

```
$ /usr/local/bin/oc adm release mirror \
-a pull-secret-update.json
--from=$UPSTREAM_REPO \
--to-release-image=$LOCAL_REG/$LOCAL_REPO:${VERSION} \
--to=$LOCAL_REG/$LOCAL_REPO
```



NOTE

Example output of the variables used to mirror the install images:

```
UPSTREAM_REPO=${RELEASE_IMAGE}
LOCAL_REG=<registry_FQDN>:<registry_port>
LOCAL_REPO='ocp4/openshift4'
```

The values of **RELEASE_IMAGE** and **VERSION** were set during the **Retrieving OpenShift Installer** step of the **Setting up the environment for an OpenShift installation** section.

2. After mirroring the registry, confirm that you can access it in your disconnected environment:

```
$ curl -k -u <user>:<password> https://registry.example.com:<registry_port>/v2/_catalog
{"repositories":["<Repo_Name>"]}
```

16.6.16. Miscellaneous issues

16.6.16.1. Addressing the runtime network not ready error

After the deployment of a cluster you might receive the following error:

```
`runtime network not ready: NetworkReady=false reason:NetworkPluginNotReady message:Network
plugin returns error: Missing CNI default network`
```

The Cluster Network Operator is responsible for deploying the networking components in response to a special object created by the installation program. It runs very early in the installation process, after the control plane (master) nodes have come up, but before the bootstrap control plane has been torn down. It can be indicative of more subtle installation program issues, such as long delays in bringing up control plane (master) nodes or issues with **apiserver** communication.

Procedure

1. Inspect the pods in the **openshift-network-operator** namespace:

```
$ oc get all -n openshift-network-operator
```

```
NAME                                READY STATUS      RESTARTS  AGE
pod/network-operator-69dfd7b577-bg89v 0/1   ContainerCreating 0      149m
```

2. On the **provisioner** node, determine that the network configuration exists:

```
$ kubectl get network.config.openshift.io cluster -oyaml
```

```
apiVersion: config.openshift.io/v1
kind: Network
metadata:
  name: cluster
spec:
  serviceNetwork:
  - 172.30.0.0/16
  clusterNetwork:
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  networkType: OVNKubernetes
```

If it does not exist, the installation program did not create it. To determine why the installation program did not create it, execute the following:

```
$ openshift-install create manifests
```

3. Check that the **network-operator** is running:

```
$ kubectl -n openshift-network-operator get pods
```


- Retrieve the logs:

```
$ kubectl -n openshift-network-operator logs -l "name=network-operator"
```

On high availability clusters with three or more control plane nodes, the Operator will perform leader election and all other Operators will sleep. For additional details, see [Troubleshooting](#).

16.6.16.2. Addressing the "No disk found with matching rootDeviceHints" error message

After you deploy a cluster, you might receive the following error message:

```
No disk found with matching rootDeviceHints
```

To address the **No disk found with matching rootDeviceHints** error message, a temporary workaround is to change the **rootDeviceHints** to **minSizeGigabytes: 300**.

After you change the **rootDeviceHints** settings, boot the CoreOS and then verify the disk information by using the following command:

```
$ udevadm info /dev/sda
```

If you are using DL360 Gen 10 servers, be aware that they have an SD-card slot that might be assigned the **/dev/sda** device name. If no SD card is present in the server, it can cause conflicts. Ensure that the SD card slot is disabled in the server's BIOS settings.

If the **minSizeGigabytes** workaround is not fulfilling the requirements, you might need to revert **rootDeviceHints** back to **/dev/sda**. This change allows ironic images to boot successfully.

An alternative approach to fixing this problem is by using the serial ID of the disk. However, be aware that finding the serial ID can be challenging and might make the configuration file less readable. If you choose this path, ensure that you gather the serial ID using the previously documented command and incorporate it into your configuration.

16.6.16.3. Cluster nodes not getting the correct IPv6 address over DHCP

If the cluster nodes are not getting the correct IPv6 address over DHCP, check the following:

- Ensure the reserved IPv6 addresses reside outside the DHCP range.
- In the IP address reservation on the DHCP server, ensure the reservation specifies the correct DHCP Unique Identifier (DUID). For example:

```
# This is a dnsmasq dhcp reservation, 'id:00:03:00:01' is the client id and '18:db:f2:8c:d5:9f' is
the MAC Address for the NIC
id:00:03:00:01:18:db:f2:8c:d5:9f,openshift-master-1,[2620:52:0:1302::6]
```

- Ensure that route announcements are working.
- Ensure that the DHCP server is listening on the required interfaces serving the IP address ranges.

16.6.16.4. Cluster nodes not getting the correct hostname over DHCP

During IPv6 deployment, cluster nodes must get their hostname over DHCP. Sometimes the **NetworkManager** does not assign the hostname immediately. A control plane (master) node might report an error such as:

```
Failed Units: 2
NetworkManager-wait-online.service
nodeip-configuration.service
```


This error indicates that the cluster node likely booted without first receiving a hostname from the DHCP server, which causes **kubelet** to boot with a **localhost.localdomain** hostname. To address the error, force the node to renew the hostname.

Procedure

1. Retrieve the **hostname**:

```
[core@master-X ~]$ hostname
```

If the hostname is **localhost**, proceed with the following steps.



NOTE

Where **X** is the control plane node number.

2. Force the cluster node to renew the DHCP lease:

```
[core@master-X ~]$ sudo nmcli con up "<bare_metal_nic>"
```

Replace **<bare_metal_nic>** with the wired connection corresponding to the **baremetal** network.

3. Check **hostname** again:

```
[core@master-X ~]$ hostname
```

4. If the hostname is still **localhost.localdomain**, restart **NetworkManager**:

```
[core@master-X ~]$ sudo systemctl restart NetworkManager
```

5. If the hostname is still **localhost.localdomain**, wait a few minutes and check again. If the hostname remains **localhost.localdomain**, repeat the previous steps.

6. Restart the **nodeip-configuration** service:

```
[core@master-X ~]$ sudo systemctl restart nodeip-configuration.service
```

This service will reconfigure the **kubelet** service with the correct hostname references.

7. Reload the unit files definition since the kubelet changed in the previous step:

```
[core@master-X ~]$ sudo systemctl daemon-reload
```

8. Restart the **kubelet** service:

```
[core@master-X ~]$ sudo systemctl restart kubelet.service
```

9. Ensure **kubelet** booted with the correct hostname:

```
[core@master-X ~]$ sudo journalctl -fu kubelet.service
```

If the cluster node is not getting the correct hostname over DHCP after the cluster is up and running, such as during a reboot, the cluster will have a pending **csr**. **Do not** approve a **csr**, or other issues might arise.

Addressing a csr

1. Get CSRs on the cluster:

```
$ oc get csr
```

2. Verify if a pending **csr** contains **Subject Name: localhost.localdomain**:

```
$ oc get csr <pending_csr> -o jsonpath='{.spec.request}' | base64 --decode | openssl req -noout -text
```

3. Remove any **csr** that contains **Subject Name: localhost.localdomain**:

```
$ oc delete csr <wrong_csr>
```

16.6.16.5. Routes do not reach endpoints

During the installation process, it is possible to encounter a Virtual Router Redundancy Protocol (VRRP) conflict. This conflict might occur if a previously used OpenShift Container Platform node that was once part of a cluster deployment using a specific cluster name is still running but not part of the current OpenShift Container Platform cluster deployment using that same cluster name. For example, a cluster was deployed using the cluster name **openshift**, deploying three control plane (master) nodes and three worker nodes. Later, a separate install uses the same cluster name **openshift**, but this redeployment only installed three control plane (master) nodes, leaving the three worker nodes from a previous deployment in an **ON** state. This might cause a Virtual Router Identifier (VRID) conflict and a VRRP conflict.

1. Get the route:

```
$ oc get route oauth-openshift
```

2. Check the service endpoint:

```
$ oc get svc oauth-openshift
```

```
NAME          TYPE          CLUSTER-IP    EXTERNAL-IP  PORT(S)  AGE
oauth-openshift ClusterIP    172.30.19.162 <none>       443/TCP  59m
```

3. Attempt to reach the service from a control plane (master) node:

```
[core@master0 ~]$ curl -k https://172.30.19.162
```

```
{
  "kind": "Status",
  "apiVersion": "v1",
  "metadata": {
  },
  "status": "Failure",
  "message": "forbidden: User \"system:anonymous\" cannot get path \"/\"",
  "reason": "Forbidden",
  "details": {
  },
  "code": 403
}
```

4. Identify the **authentication-operator** errors from the **provisioner** node:

```
$ oc logs deployment/authentication-operator -n openshift-authentication-operator
```

```
Event(v1.ObjectReference{Kind:"Deployment", Namespace:"openshift-authentication-operator", Name:"authentication-operator", UID:"225c5bd5-b368-439b-9155-5fd3c0459d98", APIVersion:"apps/v1", ResourceVersion:"", FieldPath:""}): type: 'Normal' reason: 'OperatorStatusChanged' Status for clusteroperator/authentication changed: Degraded message changed from "IngressStateEndpointsDegraded: All 2 endpoints for oauth-server are reporting"
```

Solution

1. Ensure that the cluster name for every deployment is unique, ensuring no conflict.
2. Turn off all the rogue nodes which are not part of the cluster deployment that are using the same cluster name. Otherwise, the authentication pod of the OpenShift Container Platform cluster might never start successfully.

16.6.16.6. Failed Ignition during Firstboot

During the Firstboot, the Ignition configuration may fail.

Procedure

1. Connect to the node where the Ignition configuration failed:

```
Failed Units: 1
machine-config-daemon-firstboot.service
```

2. Restart the **machine-config-daemon-firstboot** service:

```
[core@worker-X ~]$ sudo systemctl restart machine-config-daemon-firstboot.service
```

16.6.16.7. NTP out of sync

The deployment of OpenShift Container Platform clusters depends on NTP synchronized clocks among the cluster nodes. Without synchronized clocks, the deployment may fail due to clock drift if the time difference is greater than two seconds.

Procedure

1. Check for differences in the **AGE** of the cluster nodes. For example:

```
$ oc get nodes
```

```
NAME                                STATUS ROLES  AGE  VERSION
master-0.cloud.example.com  Ready  master  145m v1.29.4
master-1.cloud.example.com  Ready  master  135m v1.29.4
master-2.cloud.example.com  Ready  master  145m v1.29.4
worker-2.cloud.example.com  Ready  worker  100m v1.29.4
```

2. Check for inconsistent timing delays due to clock drift. For example:

```
$ oc get bmh -n openshift-machine-api
```

```
master-1  error registering master-1 ipmi://<out_of_band_ip>
```

```
$ sudo timedatectl
```

```
Local time: Tue 2020-03-10 18:20:02 UTC
Universal time: Tue 2020-03-10 18:20:02 UTC
RTC time: Tue 2020-03-10 18:36:53
Time zone: UTC (UTC, +0000)
System clock synchronized: no
NTP service: active
RTC in local TZ: no
```

Addressing clock drift in existing clusters

1. Create a Butane config file including the contents of the **chrony.conf** file to be delivered to the nodes. In the following example, create **99-master-chrony.bu** to add the file to the control plane nodes. You can modify the file for worker nodes or repeat this procedure for the worker role.



NOTE

See "Creating machine configs with Butane" for information about Butane.

```
variant: openshift
version: 4.16.0
metadata:
  name: 99-master-chrony
  labels:
    machineconfiguration.openshift.io/role: master
storage:
  files:
  - path: /etc/chrony.conf
    mode: 0644
    overwrite: true
  contents:
    inline: |
```

```
server <NTP_server> iburst ❶
stratumweight 0
driftfile /var/lib/chrony/drift
rtcsync
makestep 10 3
bindcmdaddress 127.0.0.1
bindcmdaddress ::1
keyfile /etc/chrony.keys
commandkey 1
generatecommandkey
noclientlog
logchange 0.5
logdir /var/log/chrony
```

❶ Replace **<NTP_server>** with the IP address of the NTP server.

2. Use Butane to generate a **MachineConfig** object file, **99-master-chrony.yaml**, containing the configuration to be delivered to the nodes:

```
$ butane 99-master-chrony.bu -o 99-master-chrony.yaml
```

3. Apply the **MachineConfig** object file:

```
$ oc apply -f 99-master-chrony.yaml
```

4. Ensure the **System clock synchronized** value is **yes**:

```
$ sudo timedatectl
```

```
Local time: Tue 2020-03-10 19:10:02 UTC
Universal time: Tue 2020-03-10 19:10:02 UTC
RTC time: Tue 2020-03-10 19:36:53
Time zone: UTC (UTC, +0000)
System clock synchronized: yes
NTP service: active
RTC in local TZ: no
```

To setup clock synchronization prior to deployment, generate the manifest files and add this file to the **openshift** directory. For example:

```
$ cp chrony-masters.yaml ~/clusterconfigs/openshift/99_masters-chrony-configuration.yaml
```

Then, continue to create the cluster.

16.6.17. Reviewing the installation

After installation, ensure the installation program deployed the nodes and pods successfully.

Procedure

1. When the OpenShift Container Platform cluster nodes are installed appropriately, the following **Ready** state is seen within the **STATUS** column:

-

```
$ oc get nodes
```

NAME	STATUS	ROLES	AGE	VERSION
master-0.example.com	Ready	master,worker	4h	v1.29.4
master-1.example.com	Ready	master,worker	4h	v1.29.4
master-2.example.com	Ready	master,worker	4h	v1.29.4


2. Confirm the installation program deployed all pods successfully. The following command removes any pods that are still running or have completed as part of the output.

```
$ oc get pods --all-namespaces | grep -iv running | grep -iv complete
```

CHAPTER 17. INSTALLING IBM CLOUD BARE METAL (CLASSIC)

17.1. PREREQUISITES

You can use installer-provisioned installation to install OpenShift Container Platform on IBM Cloud® Bare Metal (Classic) nodes. This document describes the prerequisites and procedures when installing OpenShift Container Platform on IBM Cloud® nodes.



IMPORTANT

Red Hat supports IPMI and PXE on the provisioning network only. Red Hat has not tested Red Fish, virtual media, or other complementary technologies such as Secure Boot on IBM Cloud® deployments. A provisioning network is required.


Installer-provisioned installation of OpenShift Container Platform requires:

- One node with Red Hat Enterprise Linux CoreOS (RHCOS) 8.x installed, for running the provisioner
- Three control plane nodes
- One routable network
- One provisioning network

Before starting an installer-provisioned installation of OpenShift Container Platform on IBM Cloud® Bare Metal (Classic), address the following prerequisites and requirements.

17.1.1. Setting up IBM Cloud Bare Metal (Classic) infrastructure

To deploy an OpenShift Container Platform cluster on IBM Cloud® Bare Metal (Classic) infrastructure, you must first provision the IBM Cloud® nodes.



IMPORTANT

Red Hat supports IPMI and PXE on the **provisioning** network only. Red Hat has not tested Red Fish, virtual media, or other complementary technologies such as Secure Boot on IBM Cloud® deployments. The **provisioning** network is required.

You can customize IBM Cloud® nodes using the IBM Cloud® API. When creating IBM Cloud® nodes, you must consider the following requirements.

Use one data center per cluster

All nodes in the OpenShift Container Platform cluster must run in the same IBM Cloud® data center.

Create public and private VLANs

Create all nodes with a single public VLAN and a single private VLAN.

Ensure subnets have sufficient IP addresses

IBM Cloud® public VLAN subnets use a /28 prefix by default, which provides 16 IP addresses. That is sufficient for a cluster consisting of three control plane nodes, four worker nodes, and two IP addresses for the API VIP and Ingress VIP on the **baremetal** network. For larger clusters, you might need a smaller

prefix.

IBM Cloud® private VLAN subnets use a **/26** prefix by default, which provides 64 IP addresses. IBM Cloud® Bare Metal (Classic) uses private network IP addresses to access the Baseboard Management Controller (BMC) of each node. OpenShift Container Platform creates an additional subnet for the **provisioning** network. Network traffic for the **provisioning** network subnet routes through the private VLAN. For larger clusters, you might need a smaller prefix.

Table 17.1. IP addresses per prefix

IP addresses	Prefix
32	/27
64	/26
128	/25
256	/24

Configuring NICs

OpenShift Container Platform deploys with two networks:

- **provisioning**: The **provisioning** network is a non-routable network used for provisioning the underlying operating system on each node that is a part of the OpenShift Container Platform cluster.
- **baremetal**: The **baremetal** network is a routable network. You can use any NIC order to interface with the **baremetal** network, provided it is not the NIC specified in the **provisioningNetworkInterface** configuration setting or the NIC associated to a node's **bootMACAddress** configuration setting for the **provisioning** network.

While the cluster nodes can contain more than two NICs, the installation process only focuses on the first two NICs. For example:

NIC	Network	VLAN
NIC1	provisioning	<provisioning_vlan>
NIC2	baremetal	<baremetal_vlan>

In the previous example, NIC1 on all control plane and worker nodes connects to the non-routable network (**provisioning**) that is only used for the installation of the OpenShift Container Platform cluster. NIC2 on all control plane and worker nodes connects to the routable **baremetal** network.

PXE	Boot order
NIC1 PXE-enabled provisioning network	1
NIC2 baremetal network.	2

**NOTE**

Ensure PXE is enabled on the NIC used for the **provisioning** network and is disabled on all other NICs.

Configuring canonical names

Clients access the OpenShift Container Platform cluster nodes over the **baremetal** network. Configure IBM Cloud® subdomains or subzones where the canonical name extension is the cluster name.

```
<cluster_name>.<domain>
```

For example:

```
test-cluster.example.com
```

Creating DNS entries

You must create DNS **A** record entries resolving to unused IP addresses on the public subnet for the following:

Usage	Host Name	IP
API	api.<cluster_name>.<domain>	<ip>
Ingress LB (apps)	*.apps.<cluster_name>.<domain>	<ip>


Control plane and worker nodes already have DNS entries after provisioning.

The following table provides an example of fully qualified domain names. The API and Nameserver addresses begin with canonical name extensions. The host names of the control plane and worker nodes are examples, so you can use any host naming convention you prefer.

Usage	Host Name	IP
API	api.<cluster_name>.<domain>	<ip>
Ingress LB (apps)	*.apps.<cluster_name>.<domain>	<ip>
Provisioner node	provisioner.<cluster_name>.<domain>	<ip>
Master-0	openshift-master-0.<cluster_name>.<domain>	<ip>
Master-1	openshift-master-1.<cluster_name>.<domain>	<ip>
Master-2	openshift-master-2.<cluster_name>.<domain>	<ip>

Usage	Host Name	IP
Worker-0	openshift-worker-0. <cluster_name>.<domain>	<ip>
Worker-1	openshift-worker-1. <cluster_name>.<domain>	<ip>
Worker-n	openshift-worker-n. <cluster_name>.<domain>	<ip>

OpenShift Container Platform includes functionality that uses cluster membership information to generate **A** records. This resolves the node names to their IP addresses. After the nodes are registered with the API, the cluster can disperse node information without using CoreDNS-mDNS. This eliminates the network traffic associated with multicast DNS.




IMPORTANT

After provisioning the IBM Cloud® nodes, you must create a DNS entry for the **api.<cluster_name>.<domain>** domain name on the external DNS because removing CoreDNS causes the local entry to disappear. Failure to create a DNS record for the **api.<cluster_name>.<domain>** domain name in the external DNS server prevents worker nodes from joining the cluster.

Network Time Protocol (NTP)

Each OpenShift Container Platform node in the cluster must have access to an NTP server. OpenShift Container Platform nodes use NTP to synchronize their clocks. For example, cluster nodes use SSL certificates that require validation, which might fail if the date and time between the nodes are not in sync.



IMPORTANT

Define a consistent clock date and time format in each cluster node's BIOS settings, or installation might fail.

Configure a DHCP server

IBM Cloud® Bare Metal (Classic) does not run DHCP on the public or private VLANs. After provisioning IBM Cloud® nodes, you must set up a DHCP server for the public VLAN, which corresponds to OpenShift Container Platform's **baremetal** network.



NOTE

The IP addresses allocated to each node do not need to match the IP addresses allocated by the IBM Cloud® Bare Metal (Classic) provisioning system.

See the "Configuring the public subnet" section for details.

Ensure BMC access privileges

The "Remote management" page for each node on the dashboard contains the node's intelligent platform management interface (IPMI) credentials. The default IPMI privileges prevent the user from

making certain boot target changes. You must change the privilege level to **OPERATOR** so that Ironic can make those changes.

In the `install-config.yaml` file, add the `privilegelevel` parameter to the URLs used to configure each BMC. See the "Configuring the `install-config.yaml` file" section for additional details. For example:

```
ipmi://<IP>:<port>?privilegelevel=OPERATOR
```

Alternatively, contact IBM Cloud® support and request that they increase the IPMI privileges to **ADMINISTRATOR** for each node.

Create bare metal servers

Create bare metal servers in the [IBM Cloud® dashboard](#) by navigating to **Create resource** → **Bare Metal Servers for Classic**.

Alternatively, you can create bare metal servers with the `ibmcloud` CLI utility. For example:

```
$ ibmcloud sl hardware create --hostname <SERVERNAME> \
    --domain <DOMAIN> \
    --size <SIZE> \
    --os <OS-TYPE> \
    --datacenter <DC-NAME> \
    --port-speed <SPEED> \
    --billing <BILLING>
```

See [Installing the stand-alone IBM Cloud® CLI](#) for details on installing the IBM Cloud® CLI.



NOTE

IBM Cloud® servers might take 3-5 hours to become available.

17.2. SETTING UP THE ENVIRONMENT FOR AN OPENSIFT CONTAINER PLATFORM INSTALLATION

17.2.1. Preparing the provisioner node on IBM Cloud(R) Bare Metal (Classic) infrastructure

Perform the following steps to prepare the provisioner node.

Procedure

1. Log in to the provisioner node via `ssh`.
2. Create a non-root user (`kni`) and provide that user with `sudo` privileges:

```
# useradd kni
```

```
# passwd kni
```

```
# echo "kni ALL=(root) NOPASSWD:ALL" | tee -a /etc/sudoers.d/kni
```

```
# chmod 0440 /etc/sudoers.d/kni
```

-
- 3. Create an **ssh** key for the new user:

```
# su - kni -c "ssh-keygen -f /home/kni/.ssh/id_rsa -N ""
```


- 4. Log in as the new user on the provisioner node:

```
# su - kni
```

- 5. Use Red Hat Subscription Manager to register the provisioner node:

```
$ sudo subscription-manager register --username=<user> --password=<pass> --auto-attach
```

```
$ sudo subscription-manager repos --enable=rhel-8-for-x86_64-appstream-rpms \
--enable=rhel-8-for-x86_64-baseos-rpms
```



NOTE

For more information about Red Hat Subscription Manager, see [Using and Configuring Red Hat Subscription Manager](#).

- 6. Install the following packages:

```
$ sudo dnf install -y libvirt qemu-kvm mkisofs python3-devel jq ipmitool
```

- 7. Modify the user to add the **libvirt** group to the newly created user:

```
$ sudo usermod --append --groups libvirt kni
```

- 8. Start **firewalld**:

```
$ sudo systemctl start firewalld
```

- 9. Enable **firewalld**:

```
$ sudo systemctl enable firewalld
```

- 10. Start the **http** service:

```
$ sudo firewall-cmd --zone=public --add-service=http --permanent
```

```
$ sudo firewall-cmd --reload
```

- 11. Start and enable the **libvirtd** service:

```
$ sudo systemctl enable libvirtd --now
```

- 12. Set the ID of the provisioner node:

```
$ PRVN_HOST_ID=<ID>
```

You can view the ID with the following **ibmcloud** command:

```
$ ibmcloud sl hardware list
```

13. Set the ID of the public subnet:

```
$ PUBLICSUBNETID=<ID>
```

You can view the ID with the following **ibmcloud** command:

```
$ ibmcloud sl subnet list
```

14. Set the ID of the private subnet:

```
$ PRIVSUBNETID=<ID>
```

You can view the ID with the following **ibmcloud** command:

```
$ ibmcloud sl subnet list
```

15. Set the provisioner node public IP address:

```
$ PRVN_PUB_IP=$(ibmcloud sl hardware detail $PRVN_HOST_ID --output JSON | jq  
.primaryIpAddress -r)
```

16. Set the CIDR for the public network:

```
$ PUBLICCIDR=$(ibmcloud sl subnet detail $PUBLICSUBNETID --output JSON | jq .cidr)
```

17. Set the IP address and CIDR for the public network:

```
$ PUB_IP_CIDR=$PRVN_PUB_IP/$PUBLICCIDR
```

18. Set the gateway for the public network:

```
$ PUB_GATEWAY=$(ibmcloud sl subnet detail $PUBLICSUBNETID --output JSON | jq  
.gateway -r)
```

19. Set the private IP address of the provisioner node:

```
$ PRVN_PRIV_IP=$(ibmcloud sl hardware detail $PRVN_HOST_ID --output JSON | \  
jq .primaryBackendIpAddress -r)
```

20. Set the CIDR for the private network:

```
$ PRIVCIDR=$(ibmcloud sl subnet detail $PRIVSUBNETID --output JSON | jq .cidr)
```

21. Set the IP address and CIDR for the private network:

```
$ PRIV_IP_CIDR=$PRVN_PRIV_IP/$PRIVCIDR
```

22. Set the gateway for the private network:

```
$ PRIV_GATEWAY=$(ibmcloud sl subnet detail $PRIVSUBNETID --output JSON | jq
.gateway -r)
```

23. Set up the bridges for the **baremetal** and **provisioning** networks:

```
$ sudo nohup bash -c "
  nmcli --get-values UUID con show | xargs -n 1 nmcli con delete
  nmcli connection add ifname provisioning type bridge con-name provisioning
  nmcli con add type bridge-slave ifname eth1 master provisioning
  nmcli connection add ifname baremetal type bridge con-name baremetal
  nmcli con add type bridge-slave ifname eth2 master baremetal
  nmcli connection modify baremetal ipv4.addresses $PUB_IP_CIDR ipv4.method manual
  ipv4.gateway $PUB_GATEWAY
  nmcli connection modify provisioning ipv4.addresses 172.22.0.1/24,$PRIV_IP_CIDR
  ipv4.method manual
  nmcli connection modify provisioning +ipv4.routes \"10.0.0.0/8 $PRIV_GATEWAY\"
  nmcli con down baremetal
  nmcli con up baremetal
  nmcli con down provisioning
  nmcli con up provisioning
  init 6
"
```



NOTE

For **eth1** and **eth2**, substitute the appropriate interface name, as needed.

24. If required, SSH back into the **provisioner** node:

```
# ssh kni@provisioner.<cluster-name>.<domain>
```

25. Verify the connection bridges have been properly created:

```
$ sudo nmcli con show
```

Example output

```
NAME          UUID                                TYPE  DEVICE
baremetal     4d5133a5-8351-4bb9-bfd4-3af264801530 bridge baremetal
provisioning  43942805-017f-4d7d-a2c2-7cb3324482ed bridge provisioning
virbr0       d9bca40f-eee1-410b-8879-a2d4bb0465e7 bridge virbr0
bridge-slave-eth1 76a8ed50-c7e5-4999-b4f6-6d9014dd0812 ethernet eth1
bridge-slave-eth2 f31c3353-54b7-48de-893a-02d2b34c4736 ethernet eth2
```

26. Create a **pull-secret.txt** file:

```
$ vim pull-secret.txt
```

In a web browser, navigate to [Install on Bare Metal with user-provisioned infrastructure](#). In step 1, click **Download pull secret**. Paste the contents into the **pull-secret.txt** file and save the contents in the **kni** user's home directory.

17.2.2. Configuring the public subnet

All of the OpenShift Container Platform cluster nodes must be on the public subnet. IBM Cloud® Bare Metal (Classic) does not provide a DHCP server on the subnet. Set it up separately on the provisioner node.

You must reset the BASH variables defined when preparing the provisioner node. Rebooting the provisioner node after preparing it will delete the BASH variables previously set.

Procedure

1. Install **dnsmasq**:

```
$ sudo dnf install dnsmasq
```

2. Open the **dnsmasq** configuration file:

```
$ sudo vi /etc/dnsmasq.conf
```

3. Add the following configuration to the **dnsmasq** configuration file:

```
interface=baremetal
except-interface=lo
bind-dynamic
log-dhcp

dhcp-range=<ip_addr>,<ip_addr>,<pub_cidr> 1
dhcp-option=baremetal,121,0.0.0.0/0,<pub_gateway>,<prvn_priv_ip>,<prvn_pub_ip> 2
dhcp-hostsfile=/var/lib/dnsmasq/dnsmasq.hostsfile
```

- 1 Set the DHCP range. Replace both instances of **<ip_addr>** with one unused IP address from the public subnet so that the **dhcp-range** for the **baremetal** network begins and ends with the same the IP address. Replace **<pub_cidr>** with the CIDR of the public subnet.
- 2 Set the DHCP option. Replace **<pub_gateway>** with the IP address of the gateway for the **baremetal** network. Replace **<prvn_priv_ip>** with the IP address of the provisioner node's private IP address on the **provisioning** network. Replace **<prvn_pub_ip>** with the IP address of the provisioner node's public IP address on the **baremetal** network.

To retrieve the value for **<pub_cidr>**, execute:

```
$ ibmcloud sl subnet detail <publicsubnetid> --output JSON | jq .cidr
```

Replace **<publicsubnetid>** with the ID of the public subnet.

To retrieve the value for **<pub_gateway>**, execute:

```
$ ibmcloud sl subnet detail <publicsubnetid> --output JSON | jq .gateway -r
```

Replace **<publicsubnetid>** with the ID of the public subnet.

To retrieve the value for `<prvn_priv_ip>`, execute:

```
$ ibmcloud sl hardware detail <id> --output JSON | \
jq .primaryBackendIpAddress -r
```

Replace `<id>` with the ID of the provisioner node.

To retrieve the value for `<prvn_pub_ip>`, execute:

```
$ ibmcloud sl hardware detail <id> --output JSON | jq .primaryIpAddress -r
```

Replace `<id>` with the ID of the provisioner node.

- Obtain the list of hardware for the cluster:

```
$ ibmcloud sl hardware list
```

- Obtain the MAC addresses and IP addresses for each node:

```
$ ibmcloud sl hardware detail <id> --output JSON | \
jq '.networkComponents[] | \
"\(.primaryIpAddress) \(.macAddress)" | grep -v null
```

Replace `<id>` with the ID of the node.

Example output

```
"10.196.130.144 00:e0:ed:6a:ca:b4"
"141.125.65.215 00:e0:ed:6a:ca:b5"
```

Make a note of the MAC address and IP address of the public network. Make a separate note of the MAC address of the private network, which you will use later in the `install-config.yaml` file. Repeat this procedure for each node until you have all the public MAC and IP addresses for the public **baremetal** network, and the MAC addresses of the private **provisioning** network.

- Add the MAC and IP address pair of the public **baremetal** network for each node into the `dnsmasq.hostsfile` file:

```
$ sudo vim /var/lib/dnsmasq/dnsmasq.hostsfile
```

Example input

```
00:e0:ed:6a:ca:b5,141.125.65.215,master-0
<mac>,<ip>,master-1
<mac>,<ip>,master-2
<mac>,<ip>,worker-0
<mac>,<ip>,worker-1
...
```

Replace `<mac>`,`<ip>` with the public MAC address and public IP address of the corresponding node name.

- Start **dnsmasq**:

```
$ sudo systemctl start dnsmasq
```

8. Enable **dnsmasq** so that it starts when booting the node:

```
$ sudo systemctl enable dnsmasq
```

9. Verify **dnsmasq** is running:

```
$ sudo systemctl status dnsmasq
```

Example output

```
● dnsmasq.service - DNS caching server.  
Loaded: loaded (/usr/lib/systemd/system/dnsmasq.service; enabled; vendor preset: disabled)  
Active: active (running) since Tue 2021-10-05 05:04:14 CDT; 49s ago  
Main PID: 3101 (dnsmasq)  
Tasks: 1 (limit: 204038)  
Memory: 732.0K  
CGroup: /system.slice/dnsmasq.service  
└─3101 /usr/sbin/dnsmasq -k
```

10. Open ports **53** and **67** with UDP protocol:

```
$ sudo firewall-cmd --add-port 53/udp --permanent
```

```
$ sudo firewall-cmd --add-port 67/udp --permanent
```

11. Add **provisioning** to the external zone with masquerade:

```
$ sudo firewall-cmd --change-zone=provisioning --zone=external --permanent
```

This step ensures network address translation for IPMI calls to the management subnet.

12. Reload the **firewalld** configuration:

```
$ sudo firewall-cmd --reload
```

17.2.3. Retrieving the OpenShift Container Platform installer

Use the **stable-4.x** version of the installation program and your selected architecture to deploy the generally available stable version of OpenShift Container Platform:

```
$ export VERSION=stable-4.16
```

```
$ export RELEASE_ARCH=<architecture>
```

```
$ export RELEASE_IMAGE=$(curl -s https://mirror.openshift.com/pub/openshift-  
v4/$RELEASE_ARCH/clients/ocp/$VERSION/release.txt | grep 'Pull From: quay.io' | awk -F ' ' '{print  
$3}')
```

17.2.4. Extracting the OpenShift Container Platform installer

After retrieving the installer, the next step is to extract it.

Procedure

1. Set the environment variables:

```
$ export cmd=openshift-baremetal-install
```

```
$ export pullsecret_file=~/.pull-secret.txt
```

```
$ export extract_dir=$(pwd)
```

2. Get the **oc** binary:

```
$ curl -s https://mirror.openshift.com/pub/openshift-v4/clients/ocp/$VERSION/openshift-client-linux.tar.gz | tar zxvf - oc
```

3. Extract the installer:

```
$ sudo cp oc /usr/local/bin
```

```
$ oc adm release extract --registry-config "${pullsecret_file}" --command=$cmd --to "${extract_dir}" ${RELEASE_IMAGE}
```

```
$ sudo cp openshift-baremetal-install /usr/local/bin
```

17.2.5. Configuring the install-config.yaml file

The **install-config.yaml** file requires some additional details. Most of the information is teaching the installer and the resulting cluster enough about the available IBM Cloud® Bare Metal (Classic) hardware so that it is able to fully manage it. The material difference between installing on bare metal and installing on IBM Cloud® Bare Metal (Classic) is that you must explicitly set the privilege level for IPMI in the BMC section of the **install-config.yaml** file.

Procedure

1. Configure **install-config.yaml**. Change the appropriate variables to match the environment, including **pullSecret** and **sshKey**.

```
apiVersion: v1
baseDomain: <domain>
metadata:
  name: <cluster_name>
networking:
  machineNetwork:
    - cidr: <public-cidr>
  networkType: OVNKubernetes
compute:
  - name: worker
    replicas: 2
```


```

controlPlane:
  name: master
  replicas: 3
  platform:
    baremetal: {}
platform:
  baremetal:
    apiVIP: <api_ip>
    ingressVIP: <wildcard_ip>
    provisioningNetworkInterface: <NIC1>
    provisioningNetworkCIDR: <CIDR>
  hosts:
    - name: openshift-master-0
      role: master
      bmc:
        address: ipmi://10.196.130.145?privilegelevel=OPERATOR 1
        username: root
        password: <password>
        bootMACAddress: 00:e0:ed:6a:ca:b4 2
      rootDeviceHints:
        deviceName: "/dev/sda"
    - name: openshift-worker-0
      role: worker
      bmc:
        address: ipmi://<out-of-band-ip>?privilegelevel=OPERATOR 3
        username: <user>
        password: <password>
        bootMACAddress: <NIC1_mac_address> 4
      rootDeviceHints:
        deviceName: "/dev/sda"
  pullSecret: '<pull_secret>'
  sshKey: '<ssh_pub_key>'

```

1 3 The **bmc.address** provides a **privilegelevel** configuration setting with the value set to **OPERATOR**. This is required for IBM Cloud® Bare Metal (Classic) infrastructure.

2 4 Add the MAC address of the private **provisioning** network NIC for the corresponding node.



NOTE

You can use the **ibmcloud** command-line utility to retrieve the password.

```
$ ibmcloud sl hardware detail <id> --output JSON | \
jq ""(.networkManagementIpAddress)
(.remoteManagementAccounts[0].password)""
```

Replace **<id>** with the ID of the node.

2. Create a directory to store the cluster configuration:

```
$ mkdir ~/clusterconfigs
```

- Copy the **install-config.yaml** file into the directory:

```
$ cp install-config.yaml ~/clusterconfig
```

- Ensure all bare metal nodes are powered off prior to installing the OpenShift Container Platform cluster:

```
$ ipmitool -I lanplus -U <user> -P <password> -H <management_server_ip> power off
```

- Remove old bootstrap resources if any are left over from a previous deployment attempt:

```
for i in $(sudo virsh list | tail -n +3 | grep bootstrap | awk {'print $2'});
do
  sudo virsh destroy $i;
  sudo virsh undefine $i;
  sudo virsh vol-delete $i --pool $i;
  sudo virsh vol-delete $i.ign --pool $i;
  sudo virsh pool-destroy $i;
  sudo virsh pool-undefine $i;
done
```


17.2.6. Additional install-config parameters

See the following tables for the required parameters, the **hosts** parameter, and the **bmc** parameter for the **install-config.yaml** file.

Table 17.2. Required parameters

Parameters	Default	Description
baseDomain		The domain name for the cluster. For example, example.com .
bootMode	UEFI	The boot mode for a node. Options are legacy , UEFI , and UEFISecureBoot . If bootMode is not set, Ironic sets it while inspecting the node.
bootstrapExternalStaticDNS		The static network DNS of the bootstrap node. You must set this value when deploying a cluster with static IP addresses when there is no Dynamic Host Configuration Protocol (DHCP) server on the bare-metal network. If you do not set this value, the installation program will use the value from bootstrapExternalStaticGateway , which causes problems when the IP address values of the gateway and DNS are different.
bootstrapExternalStaticIP		The static IP address for the bootstrap VM. You must set this value when deploying a cluster with static IP addresses when there is no DHCP server on the bare-metal network.
bootstrapExternalStaticGateway		The static IP address of the gateway for the bootstrap VM. You must set this value when deploying a cluster with static IP addresses when there is no DHCP server on the bare-metal network.

Parameters	Default	Description
sshKey		The sshKey configuration setting contains the key in the <code>~/.ssh/id_rsa.pub</code> file required to access the control plane nodes and compute nodes. Typically, this key is from the provisioner node.
pullSecret		The pullSecret configuration setting contains a copy of the pull secret downloaded from the Install OpenShift on Bare Metal page when preparing the provisioner node.
metadata: name:		The name to be given to the OpenShift Container Platform cluster. For example, openshift .
networking: machineNetwork: - cidr:		The public CIDR (Classless Inter-Domain Routing) of the external network. For example, 10.0.0.0/24 .
compute: - name: worker		The OpenShift Container Platform cluster requires a name be provided for compute nodes even if there are zero nodes.
compute: replicas: 2		Replicas sets the number of compute nodes in the OpenShift Container Platform cluster.
controlPlane: name: master		The OpenShift Container Platform cluster requires a name for control plane nodes.
controlPlane: replicas: 3		Replicas sets the number of control plane nodes included as part of the OpenShift Container Platform cluster.
provisioningNetwork Interface		The name of the network interface on nodes connected to the provisioning network. For OpenShift Container Platform 4.9 and later releases, use the bootMACAddress configuration setting to enable Ironic to identify the IP address of the NIC instead of using the provisioningNetworkInterface configuration setting to identify the name of the NIC.
defaultMachinePlatfo rm		The default configuration used for machine pools without a platform configuration.

Parameters	Default	Description
apiVIPs		<p>(Optional) The virtual IP address for Kubernetes API communication.</p> <p>This setting must either be provided in the install-config.yaml file as a reserved IP from the MachineNetwork or preconfigured in the DNS so that the default name resolves correctly. Use the virtual IP address and not the FQDN when adding a value to the apiVIPs configuration setting in the install-config.yaml file. The primary IP address must be from the IPv4 network when using dual stack networking. If not set, the installation program uses api.<cluster_name>.<base_domain> to derive the IP address from the DNS.</p> <div style="display: flex; align-items: flex-start;">  <div> <p>NOTE</p> <p>Before OpenShift Container Platform 4.12, the cluster installation program only accepted an IPv4 address or an IPv6 address for the apiVIP configuration setting. From OpenShift Container Platform 4.12 or later, the apiVIP configuration setting is deprecated. Instead, use a list format for the apiVIPs configuration setting to specify an IPv4 address, an IPv6 address or both IP address formats.</p> </div> </div>
disableCertificateVerification	False	redfish and redfish-virtualmedia need this parameter to manage BMC addresses. The value should be True when using a self-signed certificate for BMC addresses.


Parameters	Default	Description
ingressVIPs		<p>(Optional) The virtual IP address for ingress traffic.</p> <p>This setting must either be provided in the install-config.yaml file as a reserved IP from the MachineNetwork or preconfigured in the DNS so that the default name resolves correctly. Use the virtual IP address and not the FQDN when adding a value to the ingressVIPs configuration setting in the install-config.yaml file. The primary IP address must be from the IPv4 network when using dual stack networking. If not set, the installation program uses test.apps.<cluster_name>.<base_domain> to derive the IP address from the DNS.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p>NOTE</p> <p>Before OpenShift Container Platform 4.12, the cluster installation program only accepted an IPv4 address or an IPv6 address for the ingressVIP configuration setting. In OpenShift Container Platform 4.12 and later, the ingressVIP configuration setting is deprecated. Instead, use a list format for the ingressVIPs configuration setting to specify an IPv4 addresses, an IPv6 addresses or both IP address formats.</p> </div> </div>

Table 17.3. Optional Parameters


Parameters	Default	Description
provisioningDHCPRange	172.22.0.10,172.22.0.100	Defines the IP range for nodes on the provisioning network.
provisioningNetworkCIDR	172.22.0.0/24	The CIDR for the network to use for provisioning. This option is required when not using the default address range on the provisioning network.
clusterProvisioningIP	The third IP address of the provisioningNetworkCIDR .	The IP address within the cluster where the provisioning services run. Defaults to the third IP address of the provisioning subnet. For example, 172.22.0.3 .
bootstrapProvisioningIP	The second IP address of the provisioningNetworkCIDR .	The IP address on the bootstrap VM where the provisioning services run while the installer is deploying the control plane (master) nodes. Defaults to the second IP address of the provisioning subnet. For example, 172.22.0.2 or 2620:52:0:1307::2 .
externalBridge	baremetal	The name of the bare-metal bridge of the hypervisor attached to the bare-metal network.

Parameters	Default	Description
provisioningBridge	provisioning	The name of the provisioning bridge on the provisioner host attached to the provisioning network.
architecture		Defines the host architecture for your cluster. Valid values are amd64 or arm64 .
defaultMachinePlatform		The default configuration used for machine pools without a platform configuration.
bootstrapOSImage		A URL to override the default operating system image for the bootstrap node. The URL must contain a SHA-256 hash of the image. For example: <a href="https://mirror.openshift.com/rhcos-<version>-qemu.qcow2.gz?sha256=<uncompressed_sha256>">https://mirror.openshift.com/rhcos-<version>-qemu.qcow2.gz?sha256=<uncompressed_sha256> .
provisioningNetwork		<p>The provisioningNetwork configuration setting determines whether the cluster uses the provisioning network. If it does, the configuration setting also determines if the cluster manages the network.</p> <p>Disabled: Set this parameter to Disabled to disable the requirement for a provisioning network. When set to Disabled, you must only use virtual media based provisioning, or bring up the cluster using the assisted installer. If Disabled and using power management, BMCs must be accessible from the bare-metal network. If Disabled, you must provide two IP addresses on the bare-metal network that are used for the provisioning services.</p> <p>Managed: Set this parameter to Managed, which is the default, to fully manage the provisioning network, including DHCP, TFTP, and so on.</p> <p>Unmanaged: Set this parameter to Unmanaged to enable the provisioning network but take care of manual configuration of DHCP. Virtual media provisioning is recommended but PXE is still available if required.</p>
httpProxy		Set this parameter to the appropriate HTTP proxy used within your environment.
httpsProxy		Set this parameter to the appropriate HTTPS proxy used within your environment.
noProxy		Set this parameter to the appropriate list of exclusions for proxy usage within your environment.

Hosts

The **hosts** parameter is a list of separate bare metal assets used to build the cluster.

Table 17.4. Hosts

Name	Default	Description
name		The name of the BareMetalHost resource to associate with the details. For example, openshift-master-0 .
role		The role of the bare metal node. Either master (control plane node) or worker (compute node).
bmc		Connection details for the baseboard management controller. See the BMC addressing section for additional details.
bootMACAddress		<p>The MAC address of the NIC that the host uses for the provisioning network. Ironic retrieves the IP address using the bootMACAddress configuration setting. Then, it binds to the host.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p>NOTE</p> <p>You must provide a valid MAC address from the host if you disabled the provisioning network.</p> </div> </div>
networkConfig		Set this optional parameter to configure the network interface of a host. See "(Optional) Configuring host network interfaces" for additional details.

17.2.7. Root device hints

The **rootDeviceHints** parameter enables the installer to provision the Red Hat Enterprise Linux CoreOS (RHCOS) image to a particular device. The installer examines the devices in the order it discovers them, and compares the discovered values with the hint values. The installer uses the first discovered device that matches the hint value. The configuration can combine multiple hints, but a device must match all hints for the installer to select it.

Table 17.5. Subfields

Subfield	Description
deviceName	A string containing a Linux device name such as /dev/vda or /dev/disk/by-path/ . It is recommended to use the /dev/disk/by-path/<device_path> link to the storage location. The hint must match the actual value exactly.
hctl	A string containing a SCSI bus address like 0:0:0:0 . The hint must match the actual value exactly.

Subfield	Description
model	A string containing a vendor-specific device identifier. The hint can be a substring of the actual value.
vendor	A string containing the name of the vendor or manufacturer of the device. The hint can be a substring of the actual value.
serialNumber	A string containing the device serial number. The hint must match the actual value exactly.
minSizeGigabytes	An integer representing the minimum size of the device in gigabytes.
wwn	A string containing the unique storage identifier. The hint must match the actual value exactly.
wwnWithExtension	A string containing the unique storage identifier with the vendor extension appended. The hint must match the actual value exactly.
wwnVendorExtension	A string containing the unique vendor storage identifier. The hint must match the actual value exactly.
rotational	A boolean indicating whether the device should be a rotating disk (true) or not (false).

Example usage

```
- name: master-0
  role: master
  bmc:
    address: ipmi://10.10.0.3:6203
    username: admin
    password: redhat
  bootMACAddress: de:ad:be:ef:00:40
  rootDeviceHints:
    deviceName: "/dev/sda"
```

17.2.8. Creating the OpenShift Container Platform manifests

1. Create the OpenShift Container Platform manifests.

```
$ ./openshift-baremetal-install --dir ~/clusterconfigs create manifests
```

```
INFO Consuming Install Config from target directory
```

```
WARNING Making control-plane schedulable by setting MastersSchedulable to true for Scheduler cluster settings
WARNING Discarding the OpenShift Manifest that was provided in the target directory because its dependencies are dirty and it needs to be regenerated
```

17.2.9. Deploying the cluster via the OpenShift Container Platform installer

Run the OpenShift Container Platform installer:

```
$ ./openshift-baremetal-install --dir ~/clusterconfigs --log-level debug create cluster
```

17.2.10. Following the installation

During the deployment process, you can check the installation's overall status by issuing the **tail** command to the **.openshift_install.log** log file in the install directory folder:

```
$ tail -f /path/to/install-dir/.openshift_install.log
```

CHAPTER 18. INSTALLING ON IBM Z AND IBM LINUXONE

18.1. PREPARING TO INSTALL ON IBM Z AND IBM LINUXONE

18.1.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- Before you begin the installation process, you must clean the installation directory. This ensures that the required installation files are created and updated during the installation process.
- You provisioned [persistent storage using OpenShift Data Foundation](#) or other supported storage protocols for your cluster. To deploy a private image registry, you must set up persistent storage with **ReadWriteMany** access.
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.



NOTE

While this document refers only to IBM Z®, all information in it also applies to IBM® LinuxONE.

18.1.2. Choosing a method to install OpenShift Container Platform on IBM Z or IBM LinuxONE

The OpenShift Container Platform installation program offers the following methods for deploying a cluster on IBM Z®:

- **Interactive:** You can deploy a cluster with the web-based [Assisted Installer](#). This method requires no setup for the installer, and is ideal for connected environments like IBM Z®.
- **Local Agent-based:** You can deploy a cluster locally with the [Agent-based Installer](#). It provides many of the benefits of the Assisted Installer, but you must download and configure the Agent-based Installer first. Configuration is done with a command line interface (CLI). This approach is ideal for disconnected networks.
- **Full control:** You can deploy a cluster on [infrastructure that you prepare and maintain](#), which provides maximum customizability. You can deploy clusters in connected or disconnected environments.

Table 18.1. IBM Z(R) installation options


	Assisted Installer	Agent- based Installer	User- provisioned installatio n	Installer- provisioned installatio n
IBM Z® with z/VM	✓	✓	✓	

	Assisted Installer	Agent- based Installer	User- provisioned installation	Installer- provisioned installation
Restricted network IBM Z® with z/VM		✓	✓	
IBM Z® with RHEL KVM	✓	✓	✓	
Restricted network IBM Z® with RHEL KVM		✓	✓	
IBM Z® in an LPAR			✓	
Restricted network IBM Z® in an LPAR			✓	

For more information about the installation process, see the [Installation process](#).

18.1.2.1. User-provisioned infrastructure installation of OpenShift Container Platform on IBM Z

User-provisioned infrastructure requires the user to provision all resources required by OpenShift Container Platform.



IMPORTANT

The steps for performing a user-provisioned infrastructure installation are provided as an example only. Installing a cluster with infrastructure you provide requires knowledge of the IBM Z® platform and the installation process of OpenShift Container Platform. Use the user-provisioned infrastructure installation instructions as a guide; you are free to create the required resources through other methods.

- **Installing a cluster with z/VM on IBM Z® and IBM® LinuxONE** You can install OpenShift Container Platform with z/VM on IBM Z® or IBM® LinuxONE infrastructure that you provision.
- **Installing a cluster with z/VM on IBM Z® and IBM® LinuxONE in a restricted network** You can install OpenShift Container Platform with z/VM on IBM Z® or IBM® LinuxONE infrastructure that you provision in a restricted or disconnected network by using an internal mirror of the installation release content. You can use this method to install a cluster that does not require an active internet connection to obtain the software components. You can also use this installation method to ensure that your clusters only use container images that satisfy your organizational controls on external content.
- **Installing a cluster with RHEL KVM on IBM Z® and IBM® LinuxONE** You can install OpenShift Container Platform with KVM on IBM Z® or IBM® LinuxONE infrastructure that you provision.
- **Installing a cluster with RHEL KVM on IBM Z® and IBM® LinuxONE in a restricted network** You can install OpenShift Container Platform with RHEL KVM on IBM Z® or IBM® LinuxONE infrastructure that you provision in a restricted or disconnected network by using an internal mirror of the installation release content. You can use this method to install a cluster that does

not require an active internet connection to obtain the software components. You can also use this installation method to ensure that your clusters only use container images that satisfy your organizational controls on external content.

- **Installing a cluster in an LPAR on IBM Z® and IBM® LinuxONE** You can install OpenShift Container Platform in a logical partition (LPAR) on IBM Z® or IBM® LinuxONE infrastructure that you provision.
- **Installing a cluster in an LPAR on IBM Z® and IBM® LinuxONE in a restricted network** You can install OpenShift Container Platform in an LPAR on IBM Z® or IBM® LinuxONE infrastructure that you provision in a restricted or disconnected network by using an internal mirror of the installation release content. You can use this method to install a cluster that does not require an active internet connection to obtain the software components. You can also use this installation method to ensure that your clusters only use container images that satisfy your organizational controls on external content.

18.2. INSTALLING A CLUSTER WITH Z/VM ON IBM Z AND IBM LINUXONE

In OpenShift Container Platform version 4.16, you can install a cluster on IBM Z® or IBM® LinuxONE infrastructure that you provision.



NOTE

While this document refers only to IBM Z®, all information in it also applies to IBM® LinuxONE.



IMPORTANT

Additional considerations exist for non-bare metal platforms. Review the information in the [guidelines for deploying OpenShift Container Platform on non-tested platforms](#) before you install an OpenShift Container Platform cluster.

18.2.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- Before you begin the installation process, you must clean the installation directory. This ensures that the required installation files are created and updated during the installation process.
- You provisioned [persistent storage using OpenShift Data Foundation](#) or other supported storage protocols for your cluster. To deploy a private image registry, you must set up persistent storage with **ReadWriteMany** access.
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.



NOTE

Be sure to also review this site list if you are configuring a proxy.

18.2.2. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.



IMPORTANT

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

18.2.3. Requirements for a cluster with user-provisioned infrastructure

For a cluster that contains user-provisioned infrastructure, you must deploy all of the required machines.

This section describes the requirements for deploying OpenShift Container Platform on user-provisioned infrastructure.

18.2.3.1. Required machines for cluster installation

The smallest OpenShift Container Platform clusters require the following hosts:

Table 18.2. Minimum required hosts

Hosts	Description
One temporary bootstrap machine	The cluster requires the bootstrap machine to deploy the OpenShift Container Platform cluster on the three control plane machines. You can remove the bootstrap machine after you install the cluster.
Three control plane machines	The control plane machines run the Kubernetes and OpenShift Container Platform services that form the control plane.
At least two compute machines, which are also known as worker machines.	The workloads requested by OpenShift Container Platform users run on the compute machines.

**IMPORTANT**

To improve high availability of your cluster, distribute the control plane machines over different z/VM instances on at least two physical machines.

The bootstrap, control plane, and compute machines must use Red Hat Enterprise Linux CoreOS (RHCOS) as the operating system.

Note that RHCOS is based on Red Hat Enterprise Linux (RHEL) 9.2 and inherits all of its hardware certifications and requirements. See [Red Hat Enterprise Linux technology capabilities and limits](#) .

18.2.3.2. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 18.3. Minimum resource requirements

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)
Bootstrap	RHCOS	4	16 GB	100 GB	N/A
Control plane	RHCOS	4	16 GB	100 GB	N/A
Compute	RHCOS	2	8 GB	100 GB	N/A

1. One physical core (IFL) provides two logical cores (threads) when SMT-2 is enabled. The hypervisor can provide two or more vCPUs.

**NOTE**

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

18.2.3.3. Minimum IBM Z system environment

You can install OpenShift Container Platform version 4.16 on the following IBM® hardware:

- IBM® z16 (all models), IBM® z15 (all models), IBM® z14 (all models)
- IBM® LinuxONE 4 (all models), IBM® LinuxONE III (all models), IBM® LinuxONE Emperor II, IBM® LinuxONE Rockhopper II

Hardware requirements

- The equivalent of six Integrated Facilities for Linux (IFL), which are SMT2 enabled, for each cluster.
- At least one network connection to both connect to the **LoadBalancer** service and to serve data for traffic outside the cluster.



NOTE

You can use dedicated or shared IFLs to assign sufficient compute resources. Resource sharing is one of the key strengths of IBM Z®. However, you must adjust capacity correctly on each hypervisor layer and ensure sufficient resources for every OpenShift Container Platform cluster.



IMPORTANT

Since the overall performance of the cluster can be impacted, the LPARs that are used to set up the OpenShift Container Platform clusters must provide sufficient compute capacity. In this context, LPAR weight management, entitlements, and CPU shares on the hypervisor level play an important role.

Operating system requirements

- One instance of z/VM 7.2 or later

On your z/VM instance, set up:

- Three guest virtual machines for OpenShift Container Platform control plane machines
- Two guest virtual machines for OpenShift Container Platform compute machines
- One guest virtual machine for the temporary OpenShift Container Platform bootstrap machine

IBM Z network connectivity requirements

To install on IBM Z® under z/VM, you require a single z/VM virtual NIC in layer 2 mode. You also need:

- A direct-attached OSA or RoCE network adapter
- A z/VM VSwitch set up. For a preferred setup, use OSA link aggregation.

Disk storage

- FICON attached disk storage (DASDs). These can be z/VM minidisks, fullpack minidisks, or dedicated DASDs, all of which must be formatted as CDL, which is the default. To reach the minimum required DASD size for Red Hat Enterprise Linux CoreOS (RHCOS) installations, you need extended address volumes (EAV). If available, use HyperPAV to ensure optimal performance.

- FCP attached disk storage

Storage / Main Memory

- 16 GB for OpenShift Container Platform control plane machines
- 8 GB for OpenShift Container Platform compute machines
- 16 GB for the temporary OpenShift Container Platform bootstrap machine

18.2.3.4. Preferred IBM Z system environment

Hardware requirements

- Three LPARS that each have the equivalent of six IFLs, which are SMT2 enabled, for each cluster.
- Two network connections to both connect to the **LoadBalancer** service and to serve data for traffic outside the cluster.
- HiperSockets that are attached to a node either directly as a device or by bridging with one z/VM VSWITCH to be transparent to the z/VM guest. To directly connect HiperSockets to a node, you must set up a gateway to the external network via a RHEL 8 guest to bridge to the HiperSockets network.

Operating system requirements

- Two or three instances of z/VM 7.2 or later for high availability

On your z/VM instances, set up:

- Three guest virtual machines for OpenShift Container Platform control plane machines, one per z/VM instance.
- At least six guest virtual machines for OpenShift Container Platform compute machines, distributed across the z/VM instances.
- One guest virtual machine for the temporary OpenShift Container Platform bootstrap machine.
- To ensure the availability of integral components in an overcommitted environment, increase the priority of the control plane by using the CP command **SET SHARE**. Do the same for infrastructure nodes, if they exist. See [SET SHARE](#) in IBM® Documentation.

IBM Z network connectivity requirements

To install on IBM Z® under z/VM, you require a single z/VM virtual NIC in layer 2 mode. You also need:

- A direct-attached OSA or RoCE network adapter
- A z/VM VSwitch set up. For a preferred setup, use OSA link aggregation.

Disk storage

- FICON attached disk storage (DASDs). These can be z/VM minidisks, fullpack minidisks, or dedicated DASDs, all of which must be formatted as CDL, which is the default. To reach the minimum required DASD size for Red Hat Enterprise Linux CoreOS (RHCOS) installations, you need extended address volumes (EAV). If available, use HyperPAV to ensure optimal performance.

- FCP attached disk storage

Storage / Main Memory

- 16 GB for OpenShift Container Platform control plane machines
- 8 GB for OpenShift Container Platform compute machines
- 16 GB for the temporary OpenShift Container Platform bootstrap machine

18.2.3.5. Certificate signing requests management

Because your cluster has limited access to automatic machine management when you use infrastructure that you provision, you must provide a mechanism for approving cluster certificate signing requests (CSRs) after installation. The **kube-controller-manager** only approves the kubelet client CSRs. The **machine-approver** cannot guarantee the validity of a serving certificate that is requested by using kubelet credentials because it cannot confirm that the correct machine issued the request. You must determine and implement a method of verifying the validity of the kubelet serving certificate requests and approving them.

Additional resources

- See [Bridging a HiperSockets LAN with a z/VM Virtual Switch](#) in IBM® Documentation.
- See [Scaling HyperPAV alias devices on Linux guests on z/VM](#) for performance optimization.
- See [Topics in LPAR performance](#) for LPAR weight management and entitlements.
- [Recommended host practices for IBM Z® & IBM® LinuxONE environments](#)

18.2.3.6. Networking requirements for user-provisioned infrastructure

All the Red Hat Enterprise Linux CoreOS (RHCOS) machines require networking to be configured in **initramfs** during boot to fetch their Ignition config files.

During the initial boot, the machines require an HTTP or HTTPS server to establish a network connection to download their Ignition config files.

The machines are configured with static IP addresses. No DHCP server is required. Ensure that the machines have persistent IP addresses and hostnames.

The Kubernetes API server must be able to resolve the node names of the cluster machines. If the API servers and worker nodes are in different zones, you can configure a default DNS search zone to allow the API server to resolve the node names. Another supported approach is to always refer to hosts by their fully-qualified domain names in both the node objects and all DNS requests.

18.2.3.6.1. Network connectivity requirements

You must configure the network connectivity between machines to allow OpenShift Container Platform cluster components to communicate. Each machine must be able to resolve the hostnames of all other machines in the cluster.

This section provides details about the ports that are required.



IMPORTANT

In connected OpenShift Container Platform environments, all nodes are required to have internet access to pull images for platform containers and provide telemetry data to Red Hat.

Table 18.4. Ports used for all-machine to all-machine communications

Protocol	Port	Description
ICMP	N/A	Network reachability tests
TCP	1936	Metrics
	9000-9999	Host level services, including the node exporter on ports 9100-9101 and the Cluster Version Operator on port 9099 .
	10250-10259	The default ports that Kubernetes reserves
UDP	4789	VXLAN
	6081	Geneve
	9000-9999	Host level services, including the node exporter on ports 9100-9101 .
	500	IPsec IKE packets
	4500	IPsec NAT-T packets
	123	Network Time Protocol (NTP) on UDP port 123 If an external NTP time server is configured, you must open UDP port 123 .
TCP/UDP	30000-32767	Kubernetes node port
ESP	N/A	IPsec Encapsulating Security Payload (ESP)

Table 18.5. Ports used for all-machine to control plane communications

Protocol	Port	Description
TCP	6443	Kubernetes API

Table 18.6. Ports used for control plane machine to control plane machine communications

Protocol	Port	Description
TCP	2379-2380	etcd server and peer ports

NTP configuration for user-provisioned infrastructure

OpenShift Container Platform clusters are configured to use a public Network Time Protocol (NTP) server by default. If you want to use a local enterprise NTP server, or if your cluster is being deployed in a disconnected network, you can configure the cluster to use a specific time server. For more information, see the documentation for *Configuring chrony time service*.

Additional resources

- [Configuring chrony time service](#)

18.2.3.7. User-provisioned DNS requirements

In OpenShift Container Platform deployments, DNS name resolution is required for the following components:

- The Kubernetes API
- The OpenShift Container Platform application wildcard
- The bootstrap, control plane, and compute machines


Reverse DNS resolution is also required for the Kubernetes API, the bootstrap machine, the control plane machines, and the compute machines.

DNS A/AAAA or CNAME records are used for name resolution and PTR records are used for reverse name resolution. The reverse records are important because Red Hat Enterprise Linux CoreOS (RHCOS) uses the reverse records to set the hostnames for all the nodes, unless the hostnames are provided by DHCP. Additionally, the reverse records are used to generate the certificate signing requests (CSR) that OpenShift Container Platform needs to operate.

The following DNS records are required for a user-provisioned OpenShift Container Platform cluster and they must be in place before installation. In each record, **<cluster_name>** is the cluster name and **<base_domain>** is the base domain that you specify in the **install-config.yaml** file. A complete DNS record takes the form: **<component>.<cluster_name>.<base_domain>..**

Table 18.7. Required DNS records

Component	Record	Description
Kubernetes API	api.<cluster_name>.<base_domain>..	A DNS A/AAAA or CNAME record, and a DNS PTR record, to identify the API load balancer. These records must be resolvable by both clients external to the cluster and from all the nodes within the cluster.

Component	Record	Description
	api-int.<cluster_name>.<base_domain>.	<p>A DNS A/AAAA or CNAME record, and a DNS PTR record, to internally identify the API load balancer. These records must be resolvable from all the nodes within the cluster.</p> <div style="display: flex; align-items: flex-start;">  <div> <p>IMPORTANT</p> <p>The API server must be able to resolve the worker nodes by the hostnames that are recorded in Kubernetes. If the API server cannot resolve the node names, then proxied API calls can fail, and you cannot retrieve logs from pods.</p> </div> </div>
Routes	*.apps.<cluster_name>.<base_domain>.	<p>A wildcard DNS A/AAAA or CNAME record that refers to the application ingress load balancer. The application ingress load balancer targets the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default. These records must be resolvable by both clients external to the cluster and from all the nodes within the cluster.</p> <p>For example, console-openshift-console.apps.<cluster_name>.<base_domain> is used as a wildcard route to the OpenShift Container Platform console.</p>
Bootstrap machine	bootstrap.<cluster_name>.<base_domain>.	A DNS A/AAAA or CNAME record, and a DNS PTR record, to identify the bootstrap machine. These records must be resolvable by the nodes within the cluster.
Control plane machines	<control_plane><n>.<cluster_name>.<base_domain>.	DNS A/AAAA or CNAME records and DNS PTR records to identify each machine for the control plane nodes. These records must be resolvable by the nodes within the cluster.
Compute machines	<compute><n>.<cluster_name>.<base_domain>.	DNS A/AAAA or CNAME records and DNS PTR records to identify each machine for the worker nodes. These records must be resolvable by the nodes within the cluster.

**NOTE**

In OpenShift Container Platform 4.4 and later, you do not need to specify etcd host and SRV records in your DNS configuration.

TIP

You can use the **dig** command to verify name and reverse name resolution. See the section on *Validating DNS resolution for user-provisioned infrastructure* for detailed validation steps.

18.2.3.7.1. Example DNS configuration for user-provisioned clusters

This section provides A and PTR record configuration samples that meet the DNS requirements for deploying OpenShift Container Platform on user-provisioned infrastructure. The samples are not meant to provide advice for choosing one DNS solution over another.

In the examples, the cluster name is **ocp4** and the base domain is **example.com**.

Example DNS A record configuration for a user-provisioned cluster

The following example is a BIND zone file that shows sample A records for name resolution in a user-provisioned cluster.

Example 18.1. Sample DNS zone database

```
$TTL 1W
@ IN SOA ns1.example.com. root (
  2019070700 ; serial
  3H ; refresh (3 hours)
  30M ; retry (30 minutes)
  2W ; expiry (2 weeks)
  1W ) ; minimum (1 week)
IN NS ns1.example.com.
IN MX 10 smtp.example.com.
;
;
ns1.example.com. IN A 192.168.1.5
smtp.example.com. IN A 192.168.1.5
;
helper.example.com. IN A 192.168.1.5
helper.ocp4.example.com. IN A 192.168.1.5
;
api.ocp4.example.com. IN A 192.168.1.5 1
api-int.ocp4.example.com. IN A 192.168.1.5 2
;
*.apps.ocp4.example.com. IN A 192.168.1.5 3
;
bootstrap.ocp4.example.com. IN A 192.168.1.96 4
;
control-plane0.ocp4.example.com. IN A 192.168.1.97 5
control-plane1.ocp4.example.com. IN A 192.168.1.98 6
control-plane2.ocp4.example.com. IN A 192.168.1.99 7
;
compute0.ocp4.example.com. IN A 192.168.1.11 8
compute1.ocp4.example.com. IN A 192.168.1.7 9
;
;EOF
```

- 1 Provides name resolution for the Kubernetes API. The record refers to the IP address of the API load balancer.
- 2 Provides name resolution for the Kubernetes API. The record refers to the IP address of the API load balancer and is used for internal cluster communications.
- 3 Provides name resolution for the wildcard routes. The record refers to the IP address of the application ingress load balancer. The application ingress load balancer targets the machines

application ingress load balancer. The application ingress load balancer targets the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.



NOTE

In the example, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

- 4 Provides name resolution for the bootstrap machine.
- 5 6 7 Provides name resolution for the control plane machines.
- 8 9 Provides name resolution for the compute machines.

Example DNS PTR record configuration for a user-provisioned cluster

The following example BIND zone file shows sample PTR records for reverse name resolution in a user-provisioned cluster.

Example 18.2. Sample DNS zone database for reverse records

```
$TTL 1W
@ IN SOA ns1.example.com. root (
    2019070700 ; serial
    3H ; refresh (3 hours)
    30M ; retry (30 minutes)
    2W ; expiry (2 weeks)
    1W ) ; minimum (1 week)
IN NS ns1.example.com.
;
5.1.168.192.in-addr.arpa. IN PTR api.ocp4.example.com. 1
5.1.168.192.in-addr.arpa. IN PTR api-int.ocp4.example.com. 2
;
96.1.168.192.in-addr.arpa. IN PTR bootstrap.ocp4.example.com. 3
;
97.1.168.192.in-addr.arpa. IN PTR control-plane0.ocp4.example.com. 4
98.1.168.192.in-addr.arpa. IN PTR control-plane1.ocp4.example.com. 5
99.1.168.192.in-addr.arpa. IN PTR control-plane2.ocp4.example.com. 6
;
11.1.168.192.in-addr.arpa. IN PTR compute0.ocp4.example.com. 7
7.1.168.192.in-addr.arpa. IN PTR compute1.ocp4.example.com. 8
;
;EOF
```

- 1 Provides reverse DNS resolution for the Kubernetes API. The PTR record refers to the record name of the API load balancer.
- 2 Provides reverse DNS resolution for the Kubernetes API. The PTR record refers to the record name of the API load balancer and is used for internal cluster communications.

- 3 Provides reverse DNS resolution for the bootstrap machine.
- 4 5 6 Provides reverse DNS resolution for the control plane machines.
- 7 8 Provides reverse DNS resolution for the compute machines.

**NOTE**

A PTR record is not required for the OpenShift Container Platform application wildcard.

18.2.3.8. Load balancing requirements for user-provisioned infrastructure

Before you install OpenShift Container Platform, you must provision the API and application Ingress load balancing infrastructure. In production scenarios, you can deploy the API and application Ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

**NOTE**

If you want to deploy the API and application Ingress load balancers with a Red Hat Enterprise Linux (RHEL) instance, you must purchase the RHEL subscription separately.

The load balancing infrastructure must meet the following requirements:

1. **API load balancer.** Provides a common endpoint for users, both human and machine, to interact with and configure the platform. Configure the following conditions:
 - Layer 4 load balancing only. This can be referred to as Raw TCP or SSL Passthrough mode.
 - A stateless load balancing algorithm. The options vary based on the load balancer implementation.

**IMPORTANT**

Do not configure session persistence for an API load balancer. Configuring session persistence for a Kubernetes API server might cause performance issues from excess application traffic for your OpenShift Container Platform cluster and the Kubernetes API that runs inside the cluster.

Configure the following ports on both the front and back of the load balancers:

Table 18.8. API load balancer

Port	Back-end machines (pool members)	Internal	External	Description
6443	Bootstrap and control plane. You remove the bootstrap machine from the load balancer after the bootstrap machine initializes the cluster control plane. You must configure the /readyz endpoint for the API server health check probe.	X	X	Kubernetes API server

Port	Back-end machines (pool members)	Internal	External	Description
22623	Bootstrap and control plane. You remove the bootstrap machine from the load balancer after the bootstrap machine initializes the cluster control plane.	X		Machine config server



NOTE

The load balancer must be configured to take a maximum of 30 seconds from the time the API server turns off the **/readyz** endpoint to the removal of the API server instance from the pool. Within the time frame after **/readyz** returns an error or becomes healthy, the endpoint must have been removed or added. Probing every 5 or 10 seconds, with two successful requests to become healthy and three to become unhealthy, are well-tested values.

2. **Application Ingress load balancer.** Provides an ingress point for application traffic flowing in from outside the cluster. A working configuration for the Ingress router is required for an OpenShift Container Platform cluster.

Configure the following conditions:

- Layer 4 load balancing only. This can be referred to as Raw TCP or SSL Passthrough mode.
- A connection-based or session-based persistence is recommended, based on the options available and types of applications that will be hosted on the platform.

TIP

If the true IP address of the client can be seen by the application Ingress load balancer, enabling source IP-based session persistence can improve performance for applications that use end-to-end TLS encryption.

Configure the following ports on both the front and back of the load balancers:

Table 18.9. Application Ingress load balancer

Port	Back-end machines (pool members)	Internal	External	Description
443	The machines that run the Ingress Controller pods, compute, or worker, by default.	X	X	HTTPS traffic
80	The machines that run the Ingress Controller pods, compute, or worker, by default.	X	X	HTTP traffic

**NOTE**

If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application Ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes.

18.2.3.8.1. Example load balancer configuration for user-provisioned clusters

This section provides an example API and application Ingress load balancer configuration that meets the load balancing requirements for user-provisioned clusters. The sample is an `/etc/haproxy/haproxy.cfg` configuration for an HAProxy load balancer. The example is not meant to provide advice for choosing one load balancing solution over another.

In the example, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

**NOTE**

If you are using HAProxy as a load balancer and SELinux is set to **enforcing**, you must ensure that the HAProxy service can bind to the configured TCP port by running **setsebool -P haproxy_connect_any=1**.

Example 18.3. Sample API and application Ingress load balancer configuration

```

global
  log      127.0.0.1 local2
  pidfile  /var/run/haproxy.pid
  maxconn 4000
  daemon
defaults
  mode          http
  log           global
  option        dontlognull
  option http-server-close
  option        redispatch
  retries       3
  timeout http-request 10s
  timeout queue 1m
  timeout connect 10s
  timeout client 1m
  timeout server 1m
  timeout http-keep-alive 10s
  timeout check 10s
  maxconn      3000
listen api-server-6443 1
  bind *:6443
  mode tcp
  option httpchk GET /readyz HTTP/1.0
  option log-health-checks
  balance roundrobin
  server bootstrap bootstrap.ocp4.example.com:6443 verify none check check-ssl inter 10s fall 2
  rise 3 backup 2
  server master0 master0.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s


```

```

fall 2 rise 3
server master1 master1.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
fall 2 rise 3
server master2 master2.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
fall 2 rise 3
listen machine-config-server-22623 3
bind *:22623
mode tcp
server bootstrap bootstrap.ocp4.example.com:22623 check inter 1s backup 4
server master0 master0.ocp4.example.com:22623 check inter 1s
server master1 master1.ocp4.example.com:22623 check inter 1s
server master2 master2.ocp4.example.com:22623 check inter 1s
listen ingress-router-443 5
bind *:443
mode tcp
balance source
server compute0 compute0.ocp4.example.com:443 check inter 1s
server compute1 compute1.ocp4.example.com:443 check inter 1s
listen ingress-router-80 6
bind *:80
mode tcp
balance source
server compute0 compute0.ocp4.example.com:80 check inter 1s
server compute1 compute1.ocp4.example.com:80 check inter 1s

```

- 1 Port **6443** handles the Kubernetes API traffic and points to the control plane machines.
- 2 4 The bootstrap entries must be in place before the OpenShift Container Platform cluster installation and they must be removed after the bootstrap process is complete.
- 3 Port **22623** handles the machine config server traffic and points to the control plane machines.
- 5 Port **443** handles the HTTPS traffic and points to the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.
- 6 Port **80** handles the HTTP traffic and points to the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.



NOTE

If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application Ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes.

TIP

If you are using HAProxy as a load balancer, you can check that the **haproxy** process is listening on ports **6443**, **22623**, **443**, and **80** by running **netstat -nltpu** on the HAProxy node.

18.2.4. Preparing the user-provisioned infrastructure

Before you install OpenShift Container Platform on user-provisioned infrastructure, you must prepare the underlying infrastructure.

This section provides details about the high-level steps required to set up your cluster infrastructure in preparation for an OpenShift Container Platform installation. This includes configuring IP networking and network connectivity for your cluster nodes, preparing a web server for the Ignition files, enabling the required ports through your firewall, and setting up the required DNS and load balancing infrastructure.

After preparation, your cluster infrastructure must meet the requirements outlined in the *Requirements for a cluster with user-provisioned infrastructure* section.

Prerequisites

- You have reviewed the [OpenShift Container Platform 4.x Tested Integrations](#) page.
- You have reviewed the infrastructure requirements detailed in the *Requirements for a cluster with user-provisioned infrastructure* section.

Procedure

1. Set up static IP addresses.
2. Set up an HTTP or HTTPS server to provide Ignition files to the cluster nodes.
3. Ensure that your network infrastructure provides the required network connectivity between the cluster components. See the *Networking requirements for user-provisioned infrastructure* section for details about the requirements.
4. Configure your firewall to enable the ports required for the OpenShift Container Platform cluster components to communicate. See *Networking requirements for user-provisioned infrastructure* section for details about the ports that are required.



IMPORTANT

By default, port **1936** is accessible for an OpenShift Container Platform cluster, because each control plane node needs access to this port.

Avoid using the Ingress load balancer to expose this port, because doing so might result in the exposure of sensitive information, such as statistics and metrics, related to Ingress Controllers.

5. Setup the required DNS infrastructure for your cluster.
 - a. Configure DNS name resolution for the Kubernetes API, the application wildcard, the bootstrap machine, the control plane machines, and the compute machines.
 - b. Configure reverse DNS resolution for the Kubernetes API, the bootstrap machine, the control plane machines, and the compute machines.
See the *User-provisioned DNS requirements* section for more information about the OpenShift Container Platform DNS requirements.
6. Validate your DNS configuration.
 - a. From your installation node, run DNS lookups against the record names of the Kubernetes API, the wildcard routes, and the cluster nodes. Validate that the IP addresses in the

responses correspond to the correct components.

- b. From your installation node, run reverse DNS lookups against the IP addresses of the load balancer and the cluster nodes. Validate that the record names in the responses correspond to the correct components.

See the *Validating DNS resolution for user-provisioned infrastructure* section for detailed DNS validation steps.

7. Provision the required API and application ingress load balancing infrastructure. See the *Load balancing requirements for user-provisioned infrastructure* section for more information about the requirements.



NOTE

Some load balancing solutions require the DNS name resolution for the cluster nodes to be in place before the load balancing is initialized.

18.2.5. Validating DNS resolution for user-provisioned infrastructure

You can validate your DNS configuration before installing OpenShift Container Platform on user-provisioned infrastructure.



IMPORTANT

The validation steps detailed in this section must succeed before you install your cluster.

Prerequisites

- You have configured the required DNS records for your user-provisioned infrastructure.

Procedure

1. From your installation node, run DNS lookups against the record names of the Kubernetes API, the wildcard routes, and the cluster nodes. Validate that the IP addresses contained in the responses correspond to the correct components.
 - a. Perform a lookup against the Kubernetes API record name. Check that the result points to the IP address of the API load balancer:

```
$ dig +noall +answer @<nameserver_ip> api.<cluster_name>.<base_domain> 1
```

- 1 Replace **<nameserver_ip>** with the IP address of the nameserver, **<cluster_name>** with your cluster name, and **<base_domain>** with your base domain name.

Example output

```
api.ocp4.example.com. 604800 IN A 192.168.1.5
```

- b. Perform a lookup against the Kubernetes internal API record name. Check that the result points to the IP address of the API load balancer:

```
$ dig +noall +answer @<nameserver_ip> api-int.<cluster_name>.<base_domain>
```

Example output

```
api-int.ocp4.example.com. 604800 IN A 192.168.1.5
```

- c. Test an example `*.apps.<cluster_name>.<base_domain>` DNS wildcard lookup. All of the application wildcard lookups must resolve to the IP address of the application ingress load balancer:

```
$ dig +noall +answer @<nameserver_ip> random.apps.<cluster_name>.<base_domain>
```

Example output

```
random.apps.ocp4.example.com. 604800 IN A 192.168.1.5
```



NOTE

In the example outputs, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

You can replace **random** with another wildcard value. For example, you can query the route to the OpenShift Container Platform console:

```
$ dig +noall +answer @<nameserver_ip> console-openshift-console.apps.<cluster_name>.<base_domain>
```

Example output

```
console-openshift-console.apps.ocp4.example.com. 604800 IN A 192.168.1.5
```

- d. Run a lookup against the bootstrap DNS record name. Check that the result points to the IP address of the bootstrap node:

```
$ dig +noall +answer @<nameserver_ip> bootstrap.<cluster_name>.<base_domain>
```

Example output

```
bootstrap.ocp4.example.com. 604800 IN A 192.168.1.96
```

- e. Use this method to perform lookups against the DNS record names for the control plane and compute nodes. Check that the results correspond to the IP addresses of each node.
2. From your installation node, run reverse DNS lookups against the IP addresses of the load balancer and the cluster nodes. Validate that the record names contained in the responses correspond to the correct components.
 - a. Perform a reverse lookup against the IP address of the API load balancer. Check that the response includes the record names for the Kubernetes API and the Kubernetes internal API:


```
$ dig +noall +answer @<nameserver_ip> -x 192.168.1.5
```

Example output

```
5.1.168.192.in-addr.arpa. 604800 IN PTR api-int.ocp4.example.com. 1  
5.1.168.192.in-addr.arpa. 604800 IN PTR api.ocp4.example.com. 2
```

- 1** Provides the record name for the Kubernetes internal API.
- 2** Provides the record name for the Kubernetes API.



NOTE

A PTR record is not required for the OpenShift Container Platform application wildcard. No validation step is needed for reverse DNS resolution against the IP address of the application ingress load balancer.

- b. Perform a reverse lookup against the IP address of the bootstrap node. Check that the result points to the DNS record name of the bootstrap node:

```
$ dig +noall +answer @<nameserver_ip> -x 192.168.1.96
```

Example output

```
96.1.168.192.in-addr.arpa. 604800 IN PTR bootstrap.ocp4.example.com.
```

- c. Use this method to perform reverse lookups against the IP addresses for the control plane and compute nodes. Check that the results correspond to the DNS record names of each node.

18.2.6. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.



NOTE

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.



NOTE

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```



NOTE

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1 Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

18.2.7. Obtaining the installation program

Before you install OpenShift Container Platform, download the installation file on your provisioning machine.

Prerequisites

- You have a machine that runs Linux, for example Red Hat Enterprise Linux 8, with 500 MB of local disk space.

Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select your infrastructure provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.



IMPORTANT

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.



IMPORTANT

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.

4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#) . This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

18.2.8. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.



IMPORTANT

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**. To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.

3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.



NOTE

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your PATH.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

18.2.9. Manually creating the installation configuration file

Installing the cluster requires that you manually create the installation configuration file.

Prerequisites

- You have an SSH public key on your local machine to provide to the installation program. The key will be used for SSH authentication onto your cluster nodes for debugging and disaster recovery.
- You have obtained the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create an installation directory to store your required installation assets in:

```
$ mkdir <installation_directory>
```



IMPORTANT

You must create a directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

2. Customize the sample **install-config.yaml** file template that is provided and save it in the **<installation_directory>**.



NOTE

You must name this configuration file **install-config.yaml**.

3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the next step of the installation process. You must back it up now.

Additional resources

- [Installation configuration parameters for IBM Z®](#)

18.2.9.1. Sample install-config.yaml file for IBM Z

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.

```
apiVersion: v1
baseDomain: example.com 1
compute: 2
- hyperthreading: Enabled 3
  name: worker
  replicas: 0 4
architecture: s390x
```

```

controlPlane: 5
  hyperthreading: Enabled 6
  name: master
  replicas: 3 7
  architecture: s390x
metadata:
  name: test 8
networking:
  clusterNetwork:
  - cidr: 10.128.0.0/14 9
    hostPrefix: 23 10
  networkType: OVNKubernetes 11
  serviceNetwork: 12
  - 172.30.0.0/16
platform:
  none: {} 13
fips: false 14
pullSecret: '{"auths": ...}' 15
sshKey: 'ssh-ed25519 AAAA...' 16

```

- 1 The base domain of the cluster. All DNS records must be sub-domains of this base and include the cluster name.
- 2 5 The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.
- 3 6 Specifies whether to enable or disable simultaneous multithreading (SMT), or hyperthreading. By default, SMT is enabled to increase the performance of the cores in your machines. You can disable it by setting the parameter value to **Disabled**. If you disable SMT, you must disable it in all cluster machines; this includes both control plane and compute machines.



NOTE

Simultaneous multithreading (SMT) is enabled by default. If SMT is not available on your OpenShift Container Platform nodes, the **hyperthreading** parameter has no effect.



IMPORTANT

If you disable **hyperthreading**, whether on your OpenShift Container Platform nodes or in the **install-config.yaml** file, ensure that your capacity planning accounts for the dramatically decreased machine performance.

- 4 You must set this value to **0** when you install OpenShift Container Platform on user-provisioned infrastructure. In installer-provisioned installations, the parameter controls the number of compute machines that the cluster creates and manages for you. In user-provisioned installations, you must manually deploy the compute machines before you finish installing the cluster.

**NOTE**

If you are installing a three-node cluster, do not deploy any compute machines when you install the Red Hat Enterprise Linux CoreOS (RHCOS) machines.

- 7 The number of control plane machines that you add to the cluster. Because the cluster uses these values as the number of etcd endpoints in the cluster, the value must match the number of control plane machines that you deploy.
- 8 The cluster name that you specified in your DNS records.
- 9 A block of IP addresses from which pod IP addresses are allocated. This block must not overlap with existing physical networks. These IP addresses are used for the pod network. If you need to access the pods from an external network, you must configure load balancers and routers to manage the traffic.

**NOTE**

Class E CIDR range is reserved for a future use. To use the Class E CIDR range, you must ensure your networking environment accepts the IP addresses within the Class E CIDR range.

- 10 The subnet prefix length to assign to each individual node. For example, if **hostPrefix** is set to **23**, then each node is assigned a **/23** subnet out of the given **cidr**, which allows for 510 ($2^{(32 - 23)} - 2$) pod IP addresses. If you are required to provide access to nodes from an external network, configure load balancers and routers to manage the traffic.
- 11 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 12 The IP address pool to use for service IP addresses. You can enter only one IP address pool. This block must not overlap with existing physical networks. If you need to access the services from an external network, configure load balancers and routers to manage the traffic.
- 13 You must set the platform to **none**. You cannot provide additional platform configuration variables for IBM Z[®] infrastructure.

**IMPORTANT**

Clusters that are installed with the platform type **none** are unable to use some features, such as managing compute machines with the Machine API. This limitation applies even if the compute machines that are attached to the cluster are installed on a platform that would normally support the feature. This parameter cannot be changed after installation.

- 14 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.



IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 15 The [pull secret from Red Hat OpenShift Cluster Manager](#). This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.
- 16 The SSH public key for the **core** user in Red Hat Enterprise Linux CoreOS (RHCOS).



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

18.2.9.2. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure


1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```

apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5

```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.



NOTE

The installation program does not support the proxy **readinessEndpoints** field.



NOTE

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.



NOTE

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

18.2.9.3. Configuring a three-node cluster

Optionally, you can deploy zero compute machines in a minimal three node cluster that consists of three control plane machines only. This provides smaller, more resource efficient clusters for cluster administrators and developers to use for testing, development, and production.

In three-node OpenShift Container Platform environments, the three control plane machines are schedulable, which means that your application workloads are scheduled to run on them.

Prerequisites

- You have an existing **install-config.yaml** file.

Procedure

- Ensure that the number of compute replicas is set to **0** in your **install-config.yaml** file, as shown in the following **compute** stanza:

```
compute:
- name: worker
  platform: {}
  replicas: 0
```



NOTE

You must set the value of the **replicas** parameter for the compute machines to **0** when you install OpenShift Container Platform on user-provisioned infrastructure, regardless of the number of compute machines you are deploying. In installer-provisioned installations, the parameter controls the number of compute machines that the cluster creates and manages for you. This does not apply to user-provisioned installations, where the compute machines are deployed manually.



NOTE

The preferred resource for control plane nodes is six vCPUs and 21 GB. For three control plane nodes this is the memory + vCPU equivalent of a minimum five-node cluster. You should back the three nodes, each installed on a 120 GB disk, with three IFLs that are SMT2 enabled. The minimum tested setup is three vCPUs and 10 GB on a 120 GB disk for each control plane node.

For three-node cluster installations, follow these next steps:

- If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your

application ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes. See the *Load balancing requirements for user-provisioned infrastructure* section for more information.

- When you create the Kubernetes manifest files in the following procedure, ensure that the **mastersSchedulable** parameter in the `<installation_directory>/manifests/cluster-scheduler-02-config.yml` file is set to **true**. This enables your application workloads to run on the control plane nodes.
- Do not deploy any compute nodes when you create the Red Hat Enterprise Linux CoreOS (RHCOS) machines.

18.2.10. Cluster Network Operator configuration

The configuration for the cluster network is specified as part of the Cluster Network Operator (CNO) configuration and stored in a custom resource (CR) object that is named **cluster**. The CR specifies the fields for the **Network** API in the **operator.openshift.io** API group.

The CNO configuration inherits the following fields during cluster installation from the **Network** API in the **Network.config.openshift.io** API group:

clusterNetwork

IP address pools from which pod IP addresses are allocated.

serviceNetwork

IP address pool for services.

defaultNetwork.type

Cluster network plugin. **OVNKubernetes** is the only supported plugin during installation.

You can specify the cluster network plugin configuration for your cluster by setting the fields for the **defaultNetwork** object in the CNO object named **cluster**.

18.2.10.1. Cluster Network Operator configuration object

The fields for the Cluster Network Operator (CNO) are described in the following table:

Table 18.10. Cluster Network Operator configuration object


Field	Type	Description
metadata.name	string	The name of the CNO object. This name is always cluster .
spec.clusterNetwork	array	A list specifying the blocks of IP addresses from which pod IP addresses are allocated and the subnet prefix length assigned to each individual node in the cluster. For example: <pre>spec: clusterNetwork: - cidr: 10.128.0.0/19 hostPrefix: 23 - cidr: 10.128.32.0/19 hostPrefix: 23</pre>

Field	Type	Description
spec.serviceNetwork	array	<p>A block of IP addresses for services. The OpenShift SDN and OVN-Kubernetes network plugins support only a single IP address block for the service network. For example:</p> <pre>spec: serviceNetwork: - 172.30.0.0/14</pre> <p>You can customize this field only in the install-config.yaml file before you create the manifests. The value is read-only in the manifest file.</p>
spec.defaultNetwork	object	Configures the network plugin for the cluster network.
spec.kubeProxyConfig	object	The fields for this object specify the kube-proxy configuration. If you are using the OVN-Kubernetes cluster network plugin, the kube-proxy configuration has no effect.

defaultNetwork object configuration

The values for the **defaultNetwork** object are defined in the following table:

Table 18.11. **defaultNetwork** object

Field	Type	Description
type	string	<p>OVNKubernetes. The Red Hat OpenShift Networking network plugin is selected during installation. This value cannot be changed after cluster installation.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 1; padding-left: 10px;"> <p>NOTE</p> <p>OpenShift Container Platform uses the OVN-Kubernetes network plugin by default. OpenShift SDN is no longer available as an installation choice for new clusters.</p> </div> </div>
ovnKubernetesConfig	object	This object is only valid for the OVN-Kubernetes network plugin.

Configuration for the OVN-Kubernetes network plugin

The following table describes the configuration fields for the OVN-Kubernetes network plugin:

Table 18.12. **ovnKubernetesConfig** object


Field	Type	Description
mtu	integer	<p>The maximum transmission unit (MTU) for the Geneve (Generic Network Virtualization Encapsulation) overlay network. This is detected automatically based on the MTU of the primary network interface. You do not normally need to override the detected MTU.</p> <p>If the auto-detected value is not what you expect it to be, confirm that the MTU on the primary network interface on your nodes is correct. You cannot use this option to change the MTU value of the primary network interface on the nodes.</p> <p>If your cluster requires different MTU values for different nodes, you must set this value to 100 less than the lowest MTU value in your cluster. For example, if some nodes in your cluster have an MTU of 9001, and some have an MTU of 1500, you must set this value to 1400.</p>
genevePort	integer	The port to use for all Geneve packets. The default value is 6081 . This value cannot be changed after cluster installation.
ipsecConfig	object	Specify a configuration object for customizing the IPsec configuration.
ipv4	object	Specifies a configuration object for IPv4 settings.
ipv6	object	Specifies a configuration object for IPv6 settings.
policyAuditConfig	object	Specify a configuration object for customizing network policy audit logging. If unset, the defaults audit log settings are used.
gatewayConfig	object	<p>Optional: Specify a configuration object for customizing how egress traffic is sent to the node gateway.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p>NOTE</p> <p>While migrating egress traffic, you can expect some disruption to workloads and service traffic until the Cluster Network Operator (CNO) successfully rolls out the changes.</p> </div> </div>

Table 18.13. `ovnKubernetesConfig.ipv4` object

Field	Type	Description
-------	------	-------------

Field	Type	Description
internalTransitSwitchSubnet	string	<p>If your existing network infrastructure overlaps with the 100.88.0.0/16 IPv4 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. The subnet for the distributed transit switch that enables east-west traffic. This subnet cannot overlap with any other subnets used by OVN-Kubernetes or on the host itself. It must be large enough to accommodate one IP address per node in your cluster.</p> <p>The default value is 100.88.0.0/16.</p>
internalJoinSubnet	string	<p>If your existing network infrastructure overlaps with the 100.64.0.0/16 IPv4 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. You must ensure that the IP address range does not overlap with any other subnet used by your OpenShift Container Platform installation. The IP address range must be larger than the maximum number of nodes that can be added to the cluster. For example, if the clusterNetwork.cidr value is 10.128.0.0/14 and the clusterNetwork.hostPrefix value is /23, then the maximum number of nodes is 2⁽²³⁻¹⁴⁾=512.</p> <p>The default value is 100.64.0.0/16.</p>

Table 18.14. `ovnKubernetesConfig.ipv6` object

Field	Type	Description
internalTransitSwitchSubnet	string	<p>If your existing network infrastructure overlaps with the fd97::<64 IPv6 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. The subnet for the distributed transit switch that enables east-west traffic. This subnet cannot overlap with any other subnets used by OVN-Kubernetes or on the host itself. It must be large enough to accommodate one IP address per node in your cluster.</p> <p>The default value is fd97::<64.</p>
internalJoinSubnet	string	<p>If your existing network infrastructure overlaps with the fd98::<64 IPv6 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. You must ensure that the IP address range does not overlap with any other subnet used by your OpenShift Container Platform installation. The IP address range must be larger than the maximum number of nodes that can be added to the cluster.</p> <p>The default value is fd98::<64.</p>

Table 18.15. `policyAuditConfig` object

Field	Type	Description
rateLimit	integer	The maximum number of messages to generate every second per node. The default value is 20 messages per second.
maxFileSize	integer	The maximum size for the audit log in bytes. The default value is 50000000 or 50 MB.
maxLogFiles	integer	The maximum number of log files that are retained.
destination	string	One of the following additional audit log targets: libc The libc syslog() function of the journald process on the host. udp:<host>:<port> A syslog server. Replace <host>:<port> with the host and port of the syslog server. unix:<file> A Unix Domain Socket file specified by <file> . null Do not send the audit logs to any additional target.
syslogFacility	string	The syslog facility, such as kern , as defined by RFC5424. The default value is local0 .

Table 18.16. **gatewayConfig** object

Field	Type	Description
routingViaHost	boolean	Set this field to true to send egress traffic from pods to the host networking stack. For highly-specialized installations and applications that rely on manually configured routes in the kernel routing table, you might want to route egress traffic to the host networking stack. By default, egress traffic is processed in OVN to exit the cluster and is not affected by specialized routes in the kernel routing table. The default value is false . This field has an interaction with the Open vSwitch hardware offloading feature. If you set this field to true , you do not receive the performance benefits of the offloading because egress traffic is processed by the host networking stack.
ipForwarding	object	You can control IP forwarding for all traffic on OVN-Kubernetes managed interfaces by using the ipForwarding specification in the Network resource. Specify Restricted to only allow IP forwarding for Kubernetes related traffic. Specify Global to allow forwarding of all IP traffic. For new installations, the default is Restricted . For updates to OpenShift Container Platform 4.14 or later, the default is Global .

Field	Type	Description
ipv4	object	Optional: Specify an object to configure the internal OVN-Kubernetes masquerade address for host to service traffic for IPv4 addresses.
ipv6	object	Optional: Specify an object to configure the internal OVN-Kubernetes masquerade address for host to service traffic for IPv6 addresses.

Table 18.17. gatewayConfig.ipv4 object

Field	Type	Description
internalMasqueradeSubnet	string	The masquerade IPv4 addresses that are used internally to enable host to service traffic. The host is configured with these IP addresses as well as the shared gateway bridge interface. The default value is 169.254.169.0/29 .

Table 18.18. gatewayConfig.ipv6 object

Field	Type	Description
internalMasqueradeSubnet	string	The masquerade IPv6 addresses that are used internally to enable host to service traffic. The host is configured with these IP addresses as well as the shared gateway bridge interface. The default value is fd69::/125 .

Table 18.19. ipsecConfig object

Field	Type	Description
mode	string	Specifies the behavior of the IPsec implementation. Must be one of the following values: <ul style="list-style-type: none"> ● Disabled: IPsec is not enabled on cluster nodes. ● External: IPsec is enabled for network traffic with external hosts. ● Full: IPsec is enabled for pod traffic and network traffic with external hosts.

Example OVN-Kubernetes configuration with IPsec enabled

```
defaultNetwork:
  type: OVNKubernetes
  ovnKubernetesConfig:
    mtu: 1400
```

```
genevePort: 6081
ipsecConfig:
mode: Full
```




IMPORTANT

Using OVNKubernetes can lead to a stack exhaustion problem on IBM Power®.

kubeProxyConfig object configuration (OpenShiftSDN container network interface only)

The values for the **kubeProxyConfig** object are defined in the following table:

Table 18.20. kubeProxyConfig object

Field	Type	Description
iptablesSyncPeriod	string	<p>The refresh period for iptables rules. The default value is 30s. Valid suffixes include s, m, and h and are described in the Go time package documentation.</p> <div style="display: flex; align-items: flex-start;">  <div> <p>NOTE</p> <p>Because of performance improvements introduced in OpenShift Container Platform 4.3 and greater, adjusting the iptablesSyncPeriod parameter is no longer necessary.</p> </div> </div>
proxyArguments.iptables-min-sync-period	array	<p>The minimum duration before refreshing iptables rules. This field ensures that the refresh does not happen too frequently. Valid suffixes include s, m, and h and are described in the Go time package. The default value is:</p> <pre>kubeProxyConfig: proxyArguments: iptables-min-sync-period: - 0s</pre>

18.2.11. Creating the Kubernetes manifest and Ignition config files

Because you must modify some cluster definition files and manually start the cluster machines, you must generate the Kubernetes manifest and Ignition config files that the cluster needs to configure the machines.

The installation configuration file transforms into the Kubernetes manifests. The manifests wrap into the Ignition configuration files, which are later used to configure the cluster machines.



IMPORTANT

- The Ignition config files that the OpenShift Container Platform installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.



NOTE

The installation program that generates the manifest and Ignition files is architecture specific and can be obtained from the [client image mirror](#). The Linux version of the installation program runs on s390x only. This installer program is also available as a Mac OS version.

Prerequisites

- You obtained the OpenShift Container Platform installation program.
- You created the **install-config.yaml** installation configuration file.

Procedure

1. Change to the directory that contains the OpenShift Container Platform installation program and generate the Kubernetes manifests for the cluster:

```
$ ./openshift-install create manifests --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the installation directory that contains the **install-config.yaml** file you created.



WARNING

If you are installing a three-node cluster, skip the following step to allow the control plane nodes to be schedulable.



IMPORTANT

When you configure control plane nodes from the default unschedulable to schedulable, additional subscriptions are required. This is because control plane nodes then become compute nodes.

2. Check that the **mastersSchedulable** parameter in the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** Kubernetes manifest file is set to **false**. This setting prevents pods from being scheduled on the control plane machines:
 - a. Open the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** file.
 - b. Locate the **mastersSchedulable** parameter and ensure that it is set to **false**.
 - c. Save and exit the file.
3. To create the Ignition configuration files, run the following command from the directory that contains the installation program:

```
$ ./openshift-install create ignition-configs --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the same installation directory.

Ignition config files are created for the bootstrap, control plane, and compute nodes in the installation directory. The **kubeadmin-password** and **kubeconfig** files are created in the **./<installation_directory>/auth** directory:

```
.
├── auth
│   ├── kubeadmin-password
│   └── kubeconfig
├── bootstrap.ign
├── master.ign
├── metadata.json
└── worker.ign
```

18.2.12. Configuring NBDE with static IP in an IBM Z or IBM LinuxONE environment

Enabling NBDE disk encryption in an IBM Z® or IBM® LinuxONE environment requires additional steps, which are described in detail in this section.

Prerequisites

- You have set up the External Tang Server. See [Network-bound disk encryption](#) for instructions.
- You have installed the **butane** utility.
- You have reviewed the instructions for how to create machine configs with Butane.

Procedure

1. Create Butane configuration files for the control plane and compute nodes. The following example of a Butane configuration for a control plane node creates a file named **master-storage.bu** for disk encryption:

```
variant: openshift
version: 4.16.0
metadata:
  name: master-storage
```

```

labels:
  machineconfiguration.openshift.io/role: master
storage:
  luks:
    - clevis:
      tang:
        - thumbprint: QcPr_NHFJammnRCA3fFMVdNBwjs
          url: http://clevis.example.com:7500
      options: ❶
        - --cipher
        - aes-cbc-essiv:sha256
      device: /dev/disk/by-partlabel/root ❷
      label: luks-root
      name: root
      wipe_volume: true
    filesystems:
      - device: /dev/mapper/root
        format: xfs
        label: root
        wipe_filesystem: true
  openshift:
    fips: true ❸

```

- ❶ The cipher option is only required if FIPS mode is enabled. Omit the entry if FIPS is disabled.
- ❷ For installations on DASD-type disks, replace with **device: /dev/disk/by-label/root**.
- ❸ Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.

2. Create a customized initramfs file to boot the machine, by running the following command:

```

$ coreos-installer pxe customize \
  /root/rhcos-bootfiles/rhcos-<release>-live-initramfs.s390x.img \
  --dest-device /dev/disk/by-id/scsi-<serial_number> --dest-karg-append \
  ip=<ip_address>::<gateway_ip>:<subnet_mask>::<network_device>:none \
  --dest-karg-append nameserver=<nameserver_ip> \
  --dest-karg-append rd.neednet=1 -o \
  /root/rhcos-bootfiles/<node_name>-initramfs.s390x.img

```



NOTE

Before first boot, you must customize the initramfs for each node in the cluster, and add PXE kernel parameters.

3. Create a parameter file that includes **ignition.platform.id=metal** and **ignition.firstboot**.

Example kernel parameter file for the control plane machine:

```
rd.neednet=1 \
```

```

console=ttySCLP0 \
coreos.inst.install_dev=/dev/dasda \ 1
ignition.firstboot ignition.platform.id=metal \
coreos.live.rootfs_url=http://<http_server>/rhcos-<version>-live-rootfs.<architecture>.img \ 2
coreos.inst.ignition_url=http://<http_server>/master.ign \ 3
ip=10.19.17.2::10.19.17.1:255.255.255.0::enb0d0:none nameserver=10.19.17.1 \
zfcplib.allow_lun_scan=0 \ 4
rd.znet=qeth,0.0.bdd0,0.0.bdd1,0.0.bdd2,layer2=1 \
rd.zfcp=0.0.5677,0x600606680g7f0056,0x034F000000000000 \ 5

```

- 1 For installations on DASD-type disks, add **coreos.inst.install_dev=/dev/dasda**. Omit this value for FCP-type disks.
- 2 Specify the location of the **rootfs** artifact for the **kernel** and **initramfs** you are booting. Only HTTP and HTTPS protocols are supported.
- 3 Specify the location of the Ignition config file. Use **master.ign** or **worker.ign**. Only HTTP and HTTPS protocols are supported.
- 4 For installations on FCP-type disks, add **zfcplib.allow_lun_scan=0**. Omit this value for DASD-type disks.
- 5 For installations on DASD-type disks, replace with **rd.dasd=0.0.3490** to specify the DASD device.



NOTE

Write all options in the parameter file as a single line and make sure you have no newline characters.

Additional resources

- [Creating machine configs with Butane](#)

18.2.13. Installing RHCOS and starting the OpenShift Container Platform bootstrap process

To install OpenShift Container Platform on IBM Z® infrastructure that you provision, you must install Red Hat Enterprise Linux CoreOS (RHCOS) on z/VM guest virtual machines. When you install RHCOS, you must provide the Ignition config file that was generated by the OpenShift Container Platform installation program for the type of machine you are installing. If you have configured suitable networking, DNS, and load balancing infrastructure, the OpenShift Container Platform bootstrap process begins automatically after the RHCOS z/VM guest virtual machines have rebooted.

Complete the following steps to create the machines.

Prerequisites

- An HTTP or HTTPS server running on your provisioning machine that is accessible to the machines you create.
- If you want to enable secure boot, you have obtained the appropriate Red Hat Product Signing Key and read [Secure boot on IBM Z and IBM LinuxONE](#) in IBM documentation.

Procedure

1. Log in to Linux on your provisioning machine.
2. Obtain the Red Hat Enterprise Linux CoreOS (RHCOS) kernel, initramfs, and rootfs files from the [RHCOS image mirror](#).



IMPORTANT

The RHCOS images might not change with every release of OpenShift Container Platform. You must download images with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Only use the appropriate kernel, initramfs, and rootfs artifacts described in the following procedure.

The file names contain the OpenShift Container Platform version number. They resemble the following examples:

- kernel: **rhcos-<version>-live-kernel-<architecture>**
- initramfs: **rhcos-<version>-live-initramfs.<architecture>.img**
- rootfs: **rhcos-<version>-live-rootfs.<architecture>.img**



NOTE

The rootfs image is the same for FCP and DASD.

3. Create parameter files. The following parameters are specific for a particular virtual machine:
 - For **ip=**, specify the following seven entries:
 - i. The IP address for the machine.
 - ii. An empty string.
 - iii. The gateway.
 - iv. The netmask.
 - v. The machine host and domain name in the form **hostname.domainname**. Omit this value to let RHCOS decide.
 - vi. The network interface name. Omit this value to let RHCOS decide.
 - vii. If you use static IP addresses, specify **none**.
 - For **coreos.inst.ignition_url=**, specify the Ignition file for the machine role. Use **bootstrap.ign**, **master.ign**, or **worker.ign**. Only HTTP and HTTPS protocols are supported.
 - For **coreos.live.rootfs_url=**, specify the matching rootfs artifact for the kernel and initramfs you are booting. Only HTTP and HTTPS protocols are supported.
 - Optional: To enable secure boot, add **coreos.inst.secure_ipi**
 - For installations on DASD-type disks, complete the following tasks:

- i. For **coreos.inst.install_dev=**, specify **/dev/dasda**.
- ii. Use **rd.dasd=** to specify the DASD where RHCOS is to be installed.
- iii. Leave all other parameters unchanged.
Example parameter file, **bootstrap-0.parm**, for the bootstrap machine:

```
rd.neednet=1 \
console=ttysclp0 \
coreos.inst.install_dev=/dev/dasda \
coreos.live.rootfs_url=http://<http_server>/rhcos-<version>-live-rootfs.
<architecture>.img \ 1
coreos.inst.ignition_url=http://<http_server>/bootstrap.ign \ 2
coreos.inst.secure_ipl \ 3
ip=172.18.78.2::172.18.78.1:255.255.255.0:::none nameserver=172.18.78.1 \
rd.znet=qeth,0.0.bdf0,0.0.bdf1,0.0.bdf2,layer2=1,portno=0 \
zfcp.allow_lun_scan=0 \
rd.dasd=0.0.3490
```

- 1 Specify the location of the **rootfs** artifact for the **kernel** and **initramfs** you are booting. Only HTTP and HTTPS protocols are supported.
- 2 Specify the location of the Ignition config file. Use **bootstrap.ign**, **master.ign**, or **worker.ign**. Only HTTP and HTTPS protocols are supported.
- 3 Optional: To enable secure boot, add **coreos.inst.secure_ipl**.

Write all options in the parameter file as a single line and make sure you have no newline characters.

- For installations on FCP-type disks, complete the following tasks:
 - i. Use **rd.zfcp=<adapter>,<wwpn>,<lun>** to specify the FCP disk where RHCOS is to be installed. For multipathing repeat this step for each additional path.



NOTE

When you install with multiple paths, you must enable multipathing directly after the installation, not at a later point in time, as this can cause problems.

- ii. Set the install device as: **coreos.inst.install_dev=/dev/disk/by-id/scsi-<serial_number>**.



NOTE

If additional LUNs are configured with NPIV, FCP requires **zfcp.allow_lun_scan=0**. If you must enable **zfcp.allow_lun_scan=1** because you use a CSI driver, for example, you must configure your NPIV so that each node cannot access the boot partition of another node.

- iii. Leave all other parameters unchanged.



IMPORTANT

Additional postinstallation steps are required to fully enable multipathing. For more information, see “Enabling multipathing with kernel arguments on RHCOS” in *Postinstallation machine configuration tasks*.

The following is an example parameter file **worker-1.parm** for a compute node with multipathing:

```
rd.neednet=1 \
console=ttysclp0 \
coreos.inst.install_dev=/dev/disk/by-id/scsi-<serial_number> \
coreos.live.rootfs_url=http://<http_server>/rhcos-<version>-live-rootfs.
<architecture>.img \
coreos.inst.ignition_url=http://<http_server>/worker.ign \
ip=172.18.78.2::172.18.78.1:255.255.255.0::none nameserver=172.18.78.1 \
rd.znet=qeth,0.0.bdf0,0.0.bdf1,0.0.bdf2,layer2=1,portno=0 \
zfcp.allow_lun_scan=0 \
rd.zfcp=0.0.1987,0x50050763070bc5e3,0x4008400B00000000 \
rd.zfcp=0.0.19C7,0x50050763070bc5e3,0x4008400B00000000 \
rd.zfcp=0.0.1987,0x50050763071bc5e3,0x4008400B00000000 \
rd.zfcp=0.0.19C7,0x50050763071bc5e3,0x4008400B00000000
```

Write all options in the parameter file as a single line and make sure you have no newline characters.

4. Transfer the initramfs, kernel, parameter files, and RHCOS images to z/VM, for example with FTP. For details about how to transfer the files with FTP and boot from the virtual reader, see [Installing under Z/VM](#).
5. Punch the files to the virtual reader of the z/VM guest virtual machine that is to become your bootstrap node.
See [PUNCH](#) in IBM Documentation.

TIP

You can use the CP PUNCH command or, if you use Linux, the **vmur** command to transfer files between two z/VM guest virtual machines.

6. Log in to CMS on the bootstrap machine.
7. IPL the bootstrap machine from the reader:

```
$ ipl c
```

See [IPL](#) in IBM Documentation.

8. Repeat this procedure for the other machines in the cluster.

18.2.13.1. Advanced RHCOS installation reference

This section illustrates the networking configuration and other advanced options that allow you to modify the Red Hat Enterprise Linux CoreOS (RHCOS) manual installation process. The following tables

describe the kernel arguments and command-line options you can use with the RHCOS live installer and the **coreos-installer** command.

18.2.13.1.1. Networking and bonding options for ISO installations

If you install RHCOS from an ISO image, you can add kernel arguments manually when you boot the image to configure networking for a node. If no networking arguments are specified, DHCP is activated in the initramfs when RHCOS detects that networking is required to fetch the Ignition config file.



IMPORTANT

When adding networking arguments manually, you must also add the **rd.neednet=1** kernel argument to bring the network up in the initramfs.

The following information provides examples for configuring networking and bonding on your RHCOS nodes for ISO installations. The examples describe how to use the **ip=**, **nameserver=**, and **bond=** kernel arguments.



NOTE

Ordering is important when adding the kernel arguments: **ip=**, **nameserver=**, and then **bond=**.

The networking options are passed to the **dracut** tool during system boot. For more information about the networking options supported by **dracut**, see the [dracut.cmdline manual page](#).

The following examples are the networking options for ISO installation.

Configuring DHCP or static IP addresses

To configure an IP address, either use DHCP (**ip=dhcp**) or set an individual static IP address (**ip=<host_ip>**). If setting a static IP, you must then identify the DNS server IP address (**nameserver=<dns_ip>**) on each node. The following example sets:

- The node's IP address to **10.10.10.2**
- The gateway address to **10.10.10.254**
- The netmask to **255.255.255.0**
- The hostname to **core0.example.com**
- The DNS server address to **4.4.4.41**
- The auto-configuration value to **none**. No auto-configuration is required when IP networking is configured statically.

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none  
nameserver=4.4.4.41
```



NOTE

When you use DHCP to configure IP addressing for the RHCOS machines, the machines also obtain the DNS server information through DHCP. For DHCP-based deployments, you can define the DNS server address that is used by the RHCOS nodes through your DHCP server configuration.

Configuring an IP address without a static hostname

You can configure an IP address without assigning a static hostname. If a static hostname is not set by the user, it will be picked up and automatically set by a reverse DNS lookup. To configure an IP address without a static hostname refer to the following example:

- The node's IP address to **10.10.10.2**
- The gateway address to **10.10.10.254**
- The netmask to **255.255.255.0**
- The DNS server address to **4.4.4.41**
- The auto-configuration value to **none**. No auto-configuration is required when IP networking is configured statically.

```
ip=10.10.10.2::10.10.10.254:255.255.255.0::enp1s0:none
nameserver=4.4.4.41
```

Specifying multiple network interfaces

You can specify multiple network interfaces by setting multiple **ip=** entries.

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none
ip=10.10.10.3::10.10.10.254:255.255.255.0:core0.example.com:enp2s0:none
```

Configuring default gateway and route

Optional: You can configure routes to additional networks by setting an **rd.route=** value.



NOTE

When you configure one or multiple networks, one default gateway is required. If the additional network gateway is different from the primary network gateway, the default gateway must be the primary network gateway.

- Run the following command to configure the default gateway:

```
ip=::10.10.10.254:::
```

- Enter the following command to configure the route for the additional network:

```
rd.route=20.20.20.0/24:20.20.20.254:enp2s0
```

Disabling DHCP on a single interface

You can disable DHCP on a single interface, such as when there are two or more network interfaces and only one interface is being used. In the example, the **enp1s0** interface has a static networking configuration and DHCP is disabled for **enp2s0**, which is not used:

■

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none
ip=:::core0.example.com:enp2s0:none
```

Combining DHCP and static IP configurations

You can combine DHCP and static IP configurations on systems with multiple network interfaces, for example:

```
ip=enp1s0:dhcp
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp2s0:none
```

Configuring VLANs on individual interfaces

Optional: You can configure VLANs on individual interfaces by using the **vlan=** parameter.

- To configure a VLAN on a network interface and use a static IP address, run the following command:

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp2s0.100:none
vlan=enp2s0.100:enp2s0
```

- To configure a VLAN on a network interface and to use DHCP, run the following command:

```
ip=enp2s0.100:dhcp
vlan=enp2s0.100:enp2s0
```

Providing multiple DNS servers

You can provide multiple DNS servers by adding a **nameserver=** entry for each server, for example:

```
nameserver=1.1.1.1
nameserver=8.8.8.8
```

Bonding multiple network interfaces to a single interface

Optional: You can bond multiple network interfaces to a single interface by using the **bond=** option. Refer to the following examples:

- The syntax for configuring a bonded interface is: **bond=<name>[:<network_interfaces>][:options]**
<name> is the bonding device name (**bond0**), **<network_interfaces>** represents a comma-separated list of physical (ethernet) interfaces (**em1,em2**), and *options* is a comma-separated list of bonding options. Enter **modinfo bonding** to see available options.
- When you create a bonded interface using **bond=**, you must specify how the IP address is assigned and other information for the bonded interface.

- To configure the bonded interface to use DHCP, set the bond's IP address to **dhcp**. For example:

```
bond=bond0:em1,em2:mode=active-backup
ip=bond0:dhcp
```

- To configure the bonded interface to use a static IP address, enter the specific IP address you want and related information. For example:

```
bond=bond0:em1,em2:mode=active-backup,fail_over_mac=1
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:bond0:none
```

Always set the **fail_over_mac=1** option in active-backup mode, to avoid problems when shared OSA/RoCE cards are used.

Bonding multiple network interfaces to a single interface

Optional: You can configure VLANs on bonded interfaces by using the **vlan=** parameter and to use DHCP, for example:

```
ip=bond0.100:dhcp
bond=bond0:em1,em2:mode=active-backup
vlan=bond0.100:bond0
```

Use the following example to configure the bonded interface with a VLAN and to use a static IP address:

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:bond0.100:none
bond=bond0:em1,em2:mode=active-backup
vlan=bond0.100:bond0
```

Using network teaming

Optional: You can use a network teaming as an alternative to bonding by using the **team=** parameter:

- The syntax for configuring a team interface is: **team=name[:network_interfaces]**
name is the team device name (**team0**) and *network_interfaces* represents a comma-separated list of physical (ethernet) interfaces (**em1, em2**).



NOTE

Teaming is planned to be deprecated when RHCOS switches to an upcoming version of RHEL. For more information, see this [Red Hat Knowledgebase Article](#).

Use the following example to configure a network team:

```
team=team0:em1,em2
ip=team0:dhcp
```

18.2.14. Waiting for the bootstrap process to complete

The OpenShift Container Platform bootstrap process begins after the cluster nodes first boot into the persistent RHCOS environment that has been installed to disk. The configuration information provided through the Ignition config files is used to initialize the bootstrap process and install OpenShift Container Platform on the machines. You must wait for the bootstrap process to complete.

Prerequisites

- You have created the Ignition config files for your cluster.
- You have configured suitable network, DNS and load balancing infrastructure.
- You have obtained the installation program and generated the Ignition config files for your cluster.
- You installed RHCOS on your cluster machines and provided the Ignition config files that the OpenShift Container Platform installation program generated.
- Your machines have direct internet access or have an HTTP or HTTPS proxy available.

Procedure

1. Monitor the bootstrap process:

```
$ ./openshift-install --dir <installation_directory> wait-for bootstrap-complete \ 1
--log-level=info 2
```

- 1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.
- 2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Example output

```
INFO Waiting up to 30m0s for the Kubernetes API at https://api.test.example.com:6443...
INFO API v1.29.4 up
INFO Waiting up to 30m0s for bootstrapping to complete...
INFO It is now safe to remove the bootstrap resources
```

The command succeeds when the Kubernetes API server signals that it has been bootstrapped on the control plane machines.

2. After the bootstrap process is complete, remove the bootstrap machine from the load balancer.



IMPORTANT

You must remove the bootstrap machine from the load balancer at this point. You can also remove or reformat the bootstrap machine itself.

18.2.15. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

18.2.16. Approving the certificate signing requests for your machines

When you add machines to a cluster, two pending certificate signing requests (CSRs) are generated for each machine that you added. You must confirm that these CSRs are approved or, if necessary, approve them yourself. The client requests must be approved first, followed by the server requests.

Prerequisites

- You added machines to your cluster.

Procedure

1. Confirm that the cluster recognizes the machines:

```
$ oc get nodes
```

Example output

NAME	STATUS	ROLES	AGE	VERSION
master-0	Ready	master	63m	v1.29.4
master-1	Ready	master	63m	v1.29.4
master-2	Ready	master	64m	v1.29.4

The output lists all of the machines that you created.



NOTE

The preceding output might not include the compute nodes, also known as worker nodes, until some CSRs are approved.

2. Review the pending CSRs and ensure that you see the client requests with the **Pending** or **Approved** status for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

NAME	AGE	REQUESTOR	CONDITION
csr-mddf5	20m	system:node:master-01.example.com	Approved,Issued
csr-z5rln	16m	system:node:worker-21.example.com	Approved,Issued

3. If the CSRs were not approved, after all of the pending CSRs for the machines you added are in **Pending** status, approve the CSRs for your cluster machines:

**NOTE**

Because the CSRs rotate automatically, approve your CSRs within an hour of adding the machines to the cluster. If you do not approve them within an hour, the certificates will rotate, and more than two certificates will be present for each node. You must approve all of these certificates. After the client CSR is approved, the Kubelet creates a secondary CSR for the serving certificate, which requires manual approval. Then, subsequent serving certificate renewal requests are automatically approved by the **machine-approver** if the Kubelet requests a new certificate with identical parameters.

**NOTE**

For clusters running on platforms that are not machine API enabled, such as bare metal and other user-provisioned infrastructure, you must implement a method of automatically approving the kubelet serving certificate requests (CSRs). If a request is not approved, then the **oc exec**, **oc rsh**, and **oc logs** commands cannot succeed, because a serving certificate is required when the API server connects to the kubelet. Any operation that contacts the Kubelet endpoint requires this certificate approval to be in place. The method must watch for new CSRs, confirm that the CSR was submitted by the **node-bootstrap** service account in the **system:node** or **system:admin** groups, and confirm the identity of the node.

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{"\n"}{{end}}{{end}}' | xargs --no-run-if-empty oc adm certificate approve
```

**NOTE**

Some Operators might not become available until some CSRs are approved.

- Now that your client requests are approved, you must review the server requests for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

```
NAME      AGE   REQUESTOR                                     CONDITION
csr-bfd72 5m26s system:node:ip-10-0-50-126.us-east-2.compute.internal
Pending
csr-c57lv 5m26s system:node:ip-10-0-95-157.us-east-2.compute.internal
Pending
...
```


5. If the remaining CSRs are not approved, and are in the **Pending** status, approve the CSRs for your cluster machines:

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

1 **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{"\n"}\n{{end}}' | xargs oc adm certificate approve
```

6. After all client and server CSRs have been approved, the machines have the **Ready** status. Verify this by running the following command:

```
$ oc get nodes
```

Example output

```
NAME      STATUS  ROLES  AGE  VERSION
master-0  Ready   master 73m  v1.29.4
master-1  Ready   master 73m  v1.29.4
master-2  Ready   master 74m  v1.29.4
worker-0  Ready   worker 11m  v1.29.4
worker-1  Ready   worker 11m  v1.29.4
```



NOTE

It can take a few minutes after approval of the server CSRs for the machines to transition to the **Ready** status.

Additional information

- For more information on CSRs, see [Certificate Signing Requests](#).

18.2.17. Initial Operator configuration

After the control plane initializes, you must immediately configure some Operators so that they all become available.

Prerequisites

- Your control plane has initialized.

Procedure

- Watch the cluster components come online:

```
$ watch -n5 oc get clusteroperators
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.16.0	True	False	False	19m
baremetal	4.16.0	True	False	False	37m
cloud-credential	4.16.0	True	False	False	40m
cluster-autoscaler	4.16.0	True	False	False	37m
config-operator	4.16.0	True	False	False	38m
console	4.16.0	True	False	False	26m
csi-snapshot-controller	4.16.0	True	False	False	37m
dns	4.16.0	True	False	False	37m
etcd	4.16.0	True	False	False	36m
image-registry	4.16.0	True	False	False	31m
ingress	4.16.0	True	False	False	30m
insights	4.16.0	True	False	False	31m
kube-apiserver	4.16.0	True	False	False	26m
kube-controller-manager	4.16.0	True	False	False	36m
kube-scheduler	4.16.0	True	False	False	36m
kube-storage-version-migrator	4.16.0	True	False	False	37m
machine-api	4.16.0	True	False	False	29m
machine-approver	4.16.0	True	False	False	37m
machine-config	4.16.0	True	False	False	36m
marketplace	4.16.0	True	False	False	37m
monitoring	4.16.0	True	False	False	29m
network	4.16.0	True	False	False	38m
node-tuning	4.16.0	True	False	False	37m
openshift-apiserver	4.16.0	True	False	False	32m
openshift-controller-manager	4.16.0	True	False	False	30m
openshift-samples	4.16.0	True	False	False	32m
operator-lifecycle-manager	4.16.0	True	False	False	37m
operator-lifecycle-manager-catalog	4.16.0	True	False	False	37m
operator-lifecycle-manager-packageserver	4.16.0	True	False	False	32m
service-ca	4.16.0	True	False	False	38m
storage	4.16.0	True	False	False	37m

2. Configure the Operators that are not available.

18.2.17.1. Image registry storage configuration

The Image Registry Operator is not initially available for platforms that do not provide default storage. After installation, you must configure your registry to use storage so that the Registry Operator is made available.

Instructions are shown for configuring a persistent volume, which is required for production clusters. Where applicable, instructions are shown for configuring an empty directory as the storage location, which is available for only non-production clusters.

Additional instructions are provided for allowing the image registry to use block storage types by using the **Recreate** rollout strategy during upgrades.

18.2.17.1.1. Configuring registry storage for IBM Z

As a cluster administrator, following installation you must configure your registry to use storage.

Prerequisites

- You have access to the cluster as a user with the **cluster-admin** role.
- You have a cluster on IBM Z®.
- You have provisioned persistent storage for your cluster, such as Red Hat OpenShift Data Foundation.



IMPORTANT

OpenShift Container Platform supports **ReadWriteOnce** access for image registry storage when you have only one replica. **ReadWriteOnce** access also requires that the registry uses the **Recreate** rollout strategy. To deploy an image registry that supports high availability with two or more replicas, **ReadWriteMany** access is required.

- Must have 100Gi capacity.

Procedure

1. To configure your registry to use storage, change the **spec.storage.pvc** in the **configs.imageregistry/cluster** resource.



NOTE

When you use shared storage, review your security settings to prevent outside access.

2. Verify that you do not have a registry pod:

```
$ oc get pod -n openshift-image-registry -l docker-registry=default
```

Example output

```
No resources found in openshift-image-registry namespace
```



NOTE

If you do have a registry pod in your output, you do not need to continue with this procedure.

3. Check the registry configuration:

```
$ oc edit configs.imageregistry.operator.openshift.io
```

Example output

```
storage:
  pvc:
    claim:
```

Leave the **claim** field blank to allow the automatic creation of an **image-registry-storage** PVC.

4. Check the **clusteroperator** status:

```
$ oc get clusteroperator image-registry
```

Example output

```
NAME          VERSION          AVAILABLE  PROGRESSING  DEGRADED  SINCE
MESSAGE
image-registry 4.16             True       False        False     6h50m
```

5. Ensure that your registry is set to managed to enable building and pushing of images.

- Run:

```
$ oc edit configs.imageregistry/cluster
```

Then, change the line

```
managementState: Removed
```

to

```
managementState: Managed
```

18.2.17.1.2. Configuring storage for the image registry in non-production clusters

You must configure storage for the Image Registry Operator. For non-production clusters, you can set the image registry to an empty directory. If you do so, all images are lost if you restart the registry.

Procedure

- To set the image registry storage to an empty directory:

```
$ oc patch configs.imageregistry.operator.openshift.io cluster --type merge --patch '{"spec": {"storage":{"emptyDir":{}}}'
```



WARNING

Configure this option for only non-production clusters.

If you run this command before the Image Registry Operator initializes its components, the **oc patch** command fails with the following error:

```
Error from server (NotFound): configs.imageregistry.operator.openshift.io "cluster" not found
```

Wait a few minutes and run the command again.

18.2.18. Completing installation on user-provisioned infrastructure

After you complete the Operator configuration, you can finish installing the cluster on infrastructure that you provide.

Prerequisites

- Your control plane has initialized.
- You have completed the initial Operator configuration.

Procedure

1. Confirm that all the cluster components are online with the following command:

```
$ watch -n5 oc get clusteroperators
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.16.0	True	False	False	19m
baremetal	4.16.0	True	False	False	37m
cloud-credential	4.16.0	True	False	False	40m
cluster-autoscaler	4.16.0	True	False	False	37m
config-operator	4.16.0	True	False	False	38m
console	4.16.0	True	False	False	26m
csi-snapshot-controller	4.16.0	True	False	False	37m
dns	4.16.0	True	False	False	37m
etcd	4.16.0	True	False	False	36m
image-registry	4.16.0	True	False	False	31m
ingress	4.16.0	True	False	False	30m
insights	4.16.0	True	False	False	31m
kube-apiserver	4.16.0	True	False	False	26m
kube-controller-manager	4.16.0	True	False	False	36m
kube-scheduler	4.16.0	True	False	False	36m
kube-storage-version-migrator	4.16.0	True	False	False	37m
machine-api	4.16.0	True	False	False	29m
machine-approver	4.16.0	True	False	False	37m
machine-config	4.16.0	True	False	False	36m
marketplace	4.16.0	True	False	False	37m
monitoring	4.16.0	True	False	False	29m
network	4.16.0	True	False	False	38m
node-tuning	4.16.0	True	False	False	37m
openshift-apiserver	4.16.0	True	False	False	32m
openshift-controller-manager	4.16.0	True	False	False	30m
openshift-samples	4.16.0	True	False	False	32m
operator-lifecycle-manager	4.16.0	True	False	False	37m
operator-lifecycle-manager-catalog	4.16.0	True	False	False	37m
operator-lifecycle-manager-packageserver	4.16.0	True	False	False	32m
service-ca	4.16.0	True	False	False	38m
storage	4.16.0	True	False	False	37m

Alternatively, the following command notifies you when all of the clusters are available. It also retrieves and displays credentials:

```
$ ./openshift-install --dir <installation_directory> wait-for install-complete 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

Example output

```
INFO Waiting up to 30m0s for the cluster to initialize...
```

The command succeeds when the Cluster Version Operator finishes deploying the OpenShift Container Platform cluster from Kubernetes API server.



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

2. Confirm that the Kubernetes API server is communicating with the pods.

- a. To view a list of all pods, use the following command:

```
$ oc get pods --all-namespaces
```

Example output

```

NAMESPACE          NAME                                     READY  STATUS
RESTARTS  AGE
openshift-apiserver-operator  openshift-apiserver-operator-85cb746d55-zqhs8  1/1
Running    1    9m
openshift-apiserver          apiserver-67b9g                                1/1  Running  0
3m
openshift-apiserver          apiserver-ljcmx                                1/1  Running  0
1m
openshift-apiserver          apiserver-z25h4                                1/1  Running  0
2m
openshift-authentication-operator  authentication-operator-69d5d8bf84-vh2n8    1/1
Running    0    5m
...
```

-
- b. View the logs for a pod that is listed in the output of the previous command by using the following command:

```
$ oc logs <pod_name> -n <namespace> 1
```

- 1 Specify the pod name and namespace, as shown in the output of the previous command.

If the pod logs display, the Kubernetes API server can communicate with the cluster machines.

3. For an installation with Fibre Channel Protocol (FCP), additional steps are required to enable multipathing. Do not enable multipathing during installation. See "Enabling multipathing with kernel arguments on RHCOS" in the *Postinstallation machine configuration tasks* documentation for more information.

Verification

If you have enabled secure boot during the OpenShift Container Platform bootstrap process, the following verification steps are required:

1. Debug the node by running the following command:

```
$ oc debug node/<node_name>
chroot /host
```

2. Confirm that secure boot is enabled by running the following command:

```
$ cat /sys/firmware/ipl/secure
```

Example output

```
1 1
```

- 1 The value is **1** if secure boot is enabled and **0** if secure boot is not enabled.

18.2.19. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service
- [How to generate SOSREPORT within OpenShift4 nodes without SSH](#) .

18.2.20. Next steps

- [Enabling multipathing with kernel arguments on RHCOS](#) .
- [Customize your cluster](#).
- If necessary, you can [opt out of remote health reporting](#) .


18.3. INSTALLING A CLUSTER WITH Z/VM ON IBM Z AND IBM LINUXONE IN A RESTRICTED NETWORK

In OpenShift Container Platform version 4.16, you can install a cluster on IBM Z[®] or IBM[®] LinuxONE infrastructure that you provision in a restricted network.



NOTE

While this document refers to only IBM Z[®], all information in it also applies to IBM[®] LinuxONE.




IMPORTANT

Additional considerations exist for non-bare metal platforms. Review the information in the [guidelines for deploying OpenShift Container Platform on non-tested platforms](#) before you install an OpenShift Container Platform cluster.

18.3.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [created a mirror registry for installation in a restricted network](#) and obtained the **imageContentSources** data for your version of OpenShift Container Platform.
- Before you begin the installation process, you must move or remove any existing installation files. This ensures that the required installation files are created and updated during the installation process.



IMPORTANT

Ensure that installation steps are done from a machine with access to the installation media.

- You provisioned [persistent storage using OpenShift Data Foundation](#) or other supported storage protocols for your cluster. To deploy a private image registry, you must set up persistent storage with **ReadWriteMany** access.
- If you use a firewall and plan to use the Telemetry service, you [configured the firewall to allow the sites](#) that your cluster requires access to.

**NOTE**

Be sure to also review this site list if you are configuring a proxy.

18.3.2. About installations in restricted networks

In OpenShift Container Platform 4.16, you can perform an installation that does not require an active connection to the internet to obtain software components. Restricted network installations can be completed using installer-provisioned infrastructure or user-provisioned infrastructure, depending on the cloud platform to which you are installing the cluster.

If you choose to perform a restricted network installation on a cloud platform, you still require access to its cloud APIs. Some cloud functions, like Amazon Web Service's Route 53 DNS and IAM services, require internet access. Depending on your network, you might require less internet access for an installation on bare metal hardware, Nutanix, or on VMware vSphere.

To complete a restricted network installation, you must create a registry that mirrors the contents of the OpenShift image registry and contains the installation media. You can create this registry on a mirror host, which can access both the internet and your closed network, or by using other methods that meet your restrictions.

**IMPORTANT**

Because of the complexity of the configuration for user-provisioned installations, consider completing a standard user-provisioned infrastructure installation before you attempt a restricted network installation using user-provisioned infrastructure. Completing this test installation might make it easier to isolate and troubleshoot any issues that might arise during your installation in a restricted network.

18.3.2.1. Additional limits

Clusters in restricted networks have the following additional limitations and restrictions:

- The **ClusterVersion** status includes an **Unable to retrieve available updates** error.
- By default, you cannot use the contents of the Developer Catalog because you cannot access the required image stream tags.

18.3.3. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to obtain the images that are necessary to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.

18.3.4. Requirements for a cluster with user-provisioned infrastructure

For a cluster that contains user-provisioned infrastructure, you must deploy all of the required machines.

This section describes the requirements for deploying OpenShift Container Platform on user-provisioned infrastructure.

18.3.4.1. Required machines for cluster installation

The smallest OpenShift Container Platform clusters require the following hosts:

Table 18.21. Minimum required hosts

Hosts	Description
One temporary bootstrap machine	The cluster requires the bootstrap machine to deploy the OpenShift Container Platform cluster on the three control plane machines. You can remove the bootstrap machine after you install the cluster.
Three control plane machines	The control plane machines run the Kubernetes and OpenShift Container Platform services that form the control plane.
At least two compute machines, which are also known as worker machines.	The workloads requested by OpenShift Container Platform users run on the compute machines.



IMPORTANT

To improve high availability of your cluster, distribute the control plane machines over different z/VM instances on at least two physical machines.

The bootstrap, control plane, and compute machines must use Red Hat Enterprise Linux CoreOS (RHCOS) as the operating system.

Note that RHCOS is based on Red Hat Enterprise Linux (RHEL) 9.2 and inherits all of its hardware certifications and requirements. See [Red Hat Enterprise Linux technology capabilities and limits](#) .

18.3.4.2. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 18.22. Minimum resource requirements

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)
Bootstrap	RHCOS	4	16 GB	100 GB	N/A
Control plane	RHCOS	4	16 GB	100 GB	N/A
Compute	RHCOS	2	8 GB	100 GB	N/A

1. One physical core (IFL) provides two logical cores (threads) when SMT-2 is enabled. The hypervisor can provide two or more vCPUs.



NOTE

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

18.3.4.3. Minimum IBM Z system environment

You can install OpenShift Container Platform version 4.16 on the following IBM® hardware:

- IBM® z16 (all models), IBM® z15 (all models), IBM® z14 (all models)
- IBM® LinuxONE 4 (all models), IBM® LinuxONE III (all models), IBM® LinuxONE Emperor II, IBM® LinuxONE Rockhopper II

Hardware requirements

- The equivalent of six Integrated Facilities for Linux (IFL), which are SMT2 enabled, for each cluster.
- At least one network connection to both connect to the **LoadBalancer** service and to serve data for traffic outside the cluster.



NOTE

You can use dedicated or shared IFLs to assign sufficient compute resources. Resource sharing is one of the key strengths of IBM Z®. However, you must adjust capacity correctly on each hypervisor layer and ensure sufficient resources for every OpenShift Container Platform cluster.



IMPORTANT

Since the overall performance of the cluster can be impacted, the LPARs that are used to set up the OpenShift Container Platform clusters must provide sufficient compute capacity. In this context, LPAR weight management, entitlements, and CPU shares on the hypervisor level play an important role.

Operating system requirements

- One instance of z/VM 7.2 or later

On your z/VM instance, set up:

- Three guest virtual machines for OpenShift Container Platform control plane machines
- Two guest virtual machines for OpenShift Container Platform compute machines
- One guest virtual machine for the temporary OpenShift Container Platform bootstrap machine

IBM Z network connectivity requirements

To install on IBM Z[®] under z/VM, you require a single z/VM virtual NIC in layer 2 mode. You also need:

- A direct-attached OSA or RoCE network adapter
- A z/VM VSwitch set up. For a preferred setup, use OSA link aggregation.

Disk storage

- FICON attached disk storage (DASDs). These can be z/VM minidisks, fullpack minidisks, or dedicated DASDs, all of which must be formatted as CDL, which is the default. To reach the minimum required DASD size for Red Hat Enterprise Linux CoreOS (RHCOS) installations, you need extended address volumes (EAV). If available, use HyperPAV to ensure optimal performance.
- FCP attached disk storage

Storage / Main Memory

- 16 GB for OpenShift Container Platform control plane machines
- 8 GB for OpenShift Container Platform compute machines
- 16 GB for the temporary OpenShift Container Platform bootstrap machine

18.3.4.4. Preferred IBM Z system environment

Hardware requirements

- Three LPARS that each have the equivalent of six IFLs, which are SMT2 enabled, for each cluster.
- Two network connections to both connect to the **LoadBalancer** service and to serve data for traffic outside the cluster.
- HiperSockets that are attached to a node either directly as a device or by bridging with one z/VM VSWITCH to be transparent to the z/VM guest. To directly connect HiperSockets to a node, you must set up a gateway to the external network via a RHEL 8 guest to bridge to the

HiperSockets network.

Operating system requirements

- Two or three instances of z/VM 7.2 or later for high availability

On your z/VM instances, set up:

- Three guest virtual machines for OpenShift Container Platform control plane machines, one per z/VM instance.
- At least six guest virtual machines for OpenShift Container Platform compute machines, distributed across the z/VM instances.
- One guest virtual machine for the temporary OpenShift Container Platform bootstrap machine.
- To ensure the availability of integral components in an overcommitted environment, increase the priority of the control plane by using the CP command **SET SHARE**. Do the same for infrastructure nodes, if they exist. See [SET SHARE](#) in IBM® Documentation.

IBM Z network connectivity requirements

To install on IBM Z® under z/VM, you require a single z/VM virtual NIC in layer 2 mode. You also need:

- A direct-attached OSA or RoCE network adapter
- A z/VM VSwitch set up. For a preferred setup, use OSA link aggregation.

Disk storage

- FICON attached disk storage (DASDs). These can be z/VM minidisks, fullpack minidisks, or dedicated DASDs, all of which must be formatted as CDL, which is the default. To reach the minimum required DASD size for Red Hat Enterprise Linux CoreOS (RHCOS) installations, you need extended address volumes (EAV). If available, use HyperPAV to ensure optimal performance.
- FCP attached disk storage

Storage / Main Memory

- 16 GB for OpenShift Container Platform control plane machines
- 8 GB for OpenShift Container Platform compute machines
- 16 GB for the temporary OpenShift Container Platform bootstrap machine

18.3.4.5. Certificate signing requests management

Because your cluster has limited access to automatic machine management when you use infrastructure that you provision, you must provide a mechanism for approving cluster certificate signing requests (CSRs) after installation. The **kube-controller-manager** only approves the kubelet client CSRs. The **machine-approver** cannot guarantee the validity of a serving certificate that is requested by using kubelet credentials because it cannot confirm that the correct machine issued the request. You must determine and implement a method of verifying the validity of the kubelet serving certificate requests and approving them.

Additional resources

- See [Bridging a HiperSockets LAN with a z/VM Virtual Switch](#) in IBM® Documentation.
- See [Scaling HyperPAV alias devices on Linux guests on z/VM](#) for performance optimization.
- See [Topics in LPAR performance](#) for LPAR weight management and entitlements.
- [Recommended host practices for IBM Z® & IBM® LinuxONE environments](#)

18.3.4.6. Networking requirements for user-provisioned infrastructure

All the Red Hat Enterprise Linux CoreOS (RHCOS) machines require networking to be configured in **initramfs** during boot to fetch their Ignition config files.

During the initial boot, the machines require an IP address configuration that is set either through a DHCP server or statically by providing the required boot options. After a network connection is established, the machines download their Ignition config files from an HTTP or HTTPS server. The Ignition config files are then used to set the exact state of each machine. The Machine Config Operator completes more changes to the machines, such as the application of new certificates or keys, after installation.

It is recommended to use a DHCP server for long-term management of the cluster machines. Ensure that the DHCP server is configured to provide persistent IP addresses, DNS server information, and hostnames to the cluster machines.



NOTE

If a DHCP service is not available for your user-provisioned infrastructure, you can instead provide the IP networking configuration and the address of the DNS server to the nodes at RHCOS install time. These can be passed as boot arguments if you are installing from an ISO image. See the *Installing RHCOS and starting the OpenShift Container Platform bootstrap process* section for more information about static IP provisioning and advanced networking options.

The Kubernetes API server must be able to resolve the node names of the cluster machines. If the API servers and worker nodes are in different zones, you can configure a default DNS search zone to allow the API server to resolve the node names. Another supported approach is to always refer to hosts by their fully-qualified domain names in both the node objects and all DNS requests.

18.3.4.6.1. Setting the cluster node hostnames through DHCP

On Red Hat Enterprise Linux CoreOS (RHCOS) machines, the hostname is set through NetworkManager. By default, the machines obtain their hostname through DHCP. If the hostname is not provided by DHCP, set statically through kernel arguments, or another method, it is obtained through a reverse DNS lookup. Reverse DNS lookup occurs after the network has been initialized on a node and can take time to resolve. Other system services can start prior to this and detect the hostname as **localhost** or similar. You can avoid this by using DHCP to provide the hostname for each cluster node.

Additionally, setting the hostnames through DHCP can bypass any manual DNS record name configuration errors in environments that have a DNS split-horizon implementation.

18.3.4.6.2. Network connectivity requirements

You must configure the network connectivity between machines to allow OpenShift Container Platform cluster components to communicate. Each machine must be able to resolve the hostnames of all other machines in the cluster.

This section provides details about the ports that are required.

Table 18.23. Ports used for all-machine to all-machine communications

Protocol	Port	Description
ICMP	N/A	Network reachability tests
TCP	1936	Metrics
	9000-9999	Host level services, including the node exporter on ports 9100-9101 and the Cluster Version Operator on port 9099 .
	10250-10259	The default ports that Kubernetes reserves
UDP	4789	VXLAN
	6081	Geneve
	9000-9999	Host level services, including the node exporter on ports 9100-9101 .
	500	IPsec IKE packets
	4500	IPsec NAT-T packets
	123	Network Time Protocol (NTP) on UDP port 123 If an external NTP time server is configured, you must open UDP port 123 .
TCP/UDP	30000-32767	Kubernetes node port
ESP	N/A	IPsec Encapsulating Security Payload (ESP)

Table 18.24. Ports used for all-machine to control plane communications

Protocol	Port	Description
TCP	6443	Kubernetes API

Table 18.25. Ports used for control plane machine to control plane machine communications

Protocol	Port	Description
TCP	2379-2380	etcd server and peer ports

NTP configuration for user-provisioned infrastructure

OpenShift Container Platform clusters are configured to use a public Network Time Protocol (NTP) server by default. If you want to use a local enterprise NTP server, or if your cluster is being deployed in a disconnected network, you can configure the cluster to use a specific time server. For more information, see the documentation for *Configuring chrony time service*.

Additional resources

- [Configuring chrony time service](#)

18.3.4.7. User-provisioned DNS requirements

In OpenShift Container Platform deployments, DNS name resolution is required for the following components:

- The Kubernetes API
- The OpenShift Container Platform application wildcard
- The bootstrap, control plane, and compute machines


Reverse DNS resolution is also required for the Kubernetes API, the bootstrap machine, the control plane machines, and the compute machines.

DNS A/AAAA or CNAME records are used for name resolution and PTR records are used for reverse name resolution. The reverse records are important because Red Hat Enterprise Linux CoreOS (RHCOS) uses the reverse records to set the hostnames for all the nodes, unless the hostnames are provided by DHCP. Additionally, the reverse records are used to generate the certificate signing requests (CSR) that OpenShift Container Platform needs to operate.

The following DNS records are required for a user-provisioned OpenShift Container Platform cluster and they must be in place before installation. In each record, **<cluster_name>** is the cluster name and **<base_domain>** is the base domain that you specify in the **install-config.yaml** file. A complete DNS record takes the form: **<component>.<cluster_name>.<base_domain>..**

Table 18.26. Required DNS records

Component	Record	Description
Kubernetes API	api.<cluster_name>.<base_domain>..	A DNS A/AAAA or CNAME record, and a DNS PTR record, to identify the API load balancer. These records must be resolvable by both clients external to the cluster and from all the nodes within the cluster.

Component	Record	Description
	api-int.<cluster_name>.<base_domain>.	<p>A DNS A/AAAA or CNAME record, and a DNS PTR record, to internally identify the API load balancer. These records must be resolvable from all the nodes within the cluster.</p> <div style="display: flex; align-items: flex-start;">  <div> <p>IMPORTANT</p> <p>The API server must be able to resolve the worker nodes by the hostnames that are recorded in Kubernetes. If the API server cannot resolve the node names, then proxied API calls can fail, and you cannot retrieve logs from pods.</p> </div> </div>
Routes	*.apps.<cluster_name>.<base_domain>.	<p>A wildcard DNS A/AAAA or CNAME record that refers to the application ingress load balancer. The application ingress load balancer targets the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default. These records must be resolvable by both clients external to the cluster and from all the nodes within the cluster.</p> <p>For example, console-openshift-console.apps.<cluster_name>.<base_domain> is used as a wildcard route to the OpenShift Container Platform console.</p>
Bootstrap machine	bootstrap.<cluster_name>.<base_domain>.	A DNS A/AAAA or CNAME record, and a DNS PTR record, to identify the bootstrap machine. These records must be resolvable by the nodes within the cluster.
Control plane machines	<control_plane><n>.<cluster_name>.<base_domain>.	DNS A/AAAA or CNAME records and DNS PTR records to identify each machine for the control plane nodes. These records must be resolvable by the nodes within the cluster.
Compute machines	<compute><n>.<cluster_name>.<base_domain>.	DNS A/AAAA or CNAME records and DNS PTR records to identify each machine for the worker nodes. These records must be resolvable by the nodes within the cluster.

**NOTE**

In OpenShift Container Platform 4.4 and later, you do not need to specify etcd host and SRV records in your DNS configuration.

TIP

You can use the **dig** command to verify name and reverse name resolution. See the section on *Validating DNS resolution for user-provisioned infrastructure* for detailed validation steps.

18.3.4.7.1. Example DNS configuration for user-provisioned clusters

This section provides A and PTR record configuration samples that meet the DNS requirements for deploying OpenShift Container Platform on user-provisioned infrastructure. The samples are not meant to provide advice for choosing one DNS solution over another.

In the examples, the cluster name is **ocp4** and the base domain is **example.com**.

Example DNS A record configuration for a user-provisioned cluster

The following example is a BIND zone file that shows sample A records for name resolution in a user-provisioned cluster.

Example 18.4. Sample DNS zone database

```
$TTL 1W
@ IN SOA ns1.example.com. root (
  2019070700 ; serial
  3H ; refresh (3 hours)
  30M ; retry (30 minutes)
  2W ; expiry (2 weeks)
  1W ) ; minimum (1 week)
IN NS ns1.example.com.
IN MX 10 smtp.example.com.
;
;
ns1.example.com. IN A 192.168.1.5
smtp.example.com. IN A 192.168.1.5
;
helper.example.com. IN A 192.168.1.5
helper.ocp4.example.com. IN A 192.168.1.5
;
api.ocp4.example.com. IN A 192.168.1.5 ①
api-int.ocp4.example.com. IN A 192.168.1.5 ②
;
*.apps.ocp4.example.com. IN A 192.168.1.5 ③
;
bootstrap.ocp4.example.com. IN A 192.168.1.96 ④
;
control-plane0.ocp4.example.com. IN A 192.168.1.97 ⑤
control-plane1.ocp4.example.com. IN A 192.168.1.98 ⑥
control-plane2.ocp4.example.com. IN A 192.168.1.99 ⑦
;
compute0.ocp4.example.com. IN A 192.168.1.11 ⑧
compute1.ocp4.example.com. IN A 192.168.1.7 ⑨
;
;EOF
```

- ① Provides name resolution for the Kubernetes API. The record refers to the IP address of the API load balancer.
- ② Provides name resolution for the Kubernetes API. The record refers to the IP address of the API load balancer and is used for internal cluster communications.
- ③ Provides name resolution for the wildcard routes. The record refers to the IP address of the

**NOTE**

In the example, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

- 4 Provides name resolution for the bootstrap machine.
- 5 6 7 Provides name resolution for the control plane machines.
- 8 9 Provides name resolution for the compute machines.

Example DNS PTR record configuration for a user-provisioned cluster

The following example BIND zone file shows sample PTR records for reverse name resolution in a user-provisioned cluster.

Example 18.5. Sample DNS zone database for reverse records

```
$TTL 1W
@ IN SOA ns1.example.com. root (
    2019070700 ; serial
    3H ; refresh (3 hours)
    30M ; retry (30 minutes)
    2W ; expiry (2 weeks)
    1W ) ; minimum (1 week)
IN NS ns1.example.com.
;
5.1.168.192.in-addr.arpa. IN PTR api.ocp4.example.com. 1
5.1.168.192.in-addr.arpa. IN PTR api-int.ocp4.example.com. 2
;
96.1.168.192.in-addr.arpa. IN PTR bootstrap.ocp4.example.com. 3
;
97.1.168.192.in-addr.arpa. IN PTR control-plane0.ocp4.example.com. 4
98.1.168.192.in-addr.arpa. IN PTR control-plane1.ocp4.example.com. 5
99.1.168.192.in-addr.arpa. IN PTR control-plane2.ocp4.example.com. 6
;
11.1.168.192.in-addr.arpa. IN PTR compute0.ocp4.example.com. 7
7.1.168.192.in-addr.arpa. IN PTR compute1.ocp4.example.com. 8
;
;EOF
```

- 1 Provides reverse DNS resolution for the Kubernetes API. The PTR record refers to the record name of the API load balancer.
- 2 Provides reverse DNS resolution for the Kubernetes API. The PTR record refers to the record name of the API load balancer and is used for internal cluster communications.
- 3 Provides reverse DNS resolution for the bootstrap machine.
- 4 5 6 Provides reverse DNS resolution for the control plane machines.

7 8 Provides reverse DNS resolution for the compute machines.



NOTE

A PTR record is not required for the OpenShift Container Platform application wildcard.

18.3.4.8. Load balancing requirements for user-provisioned infrastructure

Before you install OpenShift Container Platform, you must provision the API and application Ingress load balancing infrastructure. In production scenarios, you can deploy the API and application Ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.



NOTE

If you want to deploy the API and application Ingress load balancers with a Red Hat Enterprise Linux (RHEL) instance, you must purchase the RHEL subscription separately.

The load balancing infrastructure must meet the following requirements:

1. **API load balancer.** Provides a common endpoint for users, both human and machine, to interact with and configure the platform. Configure the following conditions:
 - Layer 4 load balancing only. This can be referred to as Raw TCP or SSL Passthrough mode.
 - A stateless load balancing algorithm. The options vary based on the load balancer implementation.



IMPORTANT

Do not configure session persistence for an API load balancer. Configuring session persistence for a Kubernetes API server might cause performance issues from excess application traffic for your OpenShift Container Platform cluster and the Kubernetes API that runs inside the cluster.

Configure the following ports on both the front and back of the load balancers:

Table 18.27. API load balancer

Port	Back-end machines (pool members)	Internal	External	Description
6443	Bootstrap and control plane. You remove the bootstrap machine from the load balancer after the bootstrap machine initializes the cluster control plane. You must configure the /readyz endpoint for the API server health check probe.	X	X	Kubernetes API server

Port	Back-end machines (pool members)	Internal	External	Description
22623	Bootstrap and control plane. You remove the bootstrap machine from the load balancer after the bootstrap machine initializes the cluster control plane.	X		Machine config server



NOTE

The load balancer must be configured to take a maximum of 30 seconds from the time the API server turns off the `/readyz` endpoint to the removal of the API server instance from the pool. Within the time frame after `/readyz` returns an error or becomes healthy, the endpoint must have been removed or added. Probing every 5 or 10 seconds, with two successful requests to become healthy and three to become unhealthy, are well-tested values.

2. **Application Ingress load balancer.** Provides an ingress point for application traffic flowing in from outside the cluster. A working configuration for the Ingress router is required for an OpenShift Container Platform cluster.

Configure the following conditions:

- Layer 4 load balancing only. This can be referred to as Raw TCP or SSL Passthrough mode.
- A connection-based or session-based persistence is recommended, based on the options available and types of applications that will be hosted on the platform.

TIP

If the true IP address of the client can be seen by the application Ingress load balancer, enabling source IP-based session persistence can improve performance for applications that use end-to-end TLS encryption.

Configure the following ports on both the front and back of the load balancers:

Table 18.28. Application Ingress load balancer

Port	Back-end machines (pool members)	Internal	External	Description
443	The machines that run the Ingress Controller pods, compute, or worker, by default.	X	X	HTTPS traffic
80	The machines that run the Ingress Controller pods, compute, or worker, by default.	X	X	HTTP traffic

**NOTE**

If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application Ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes.

18.3.4.8.1. Example load balancer configuration for user-provisioned clusters

This section provides an example API and application Ingress load balancer configuration that meets the load balancing requirements for user-provisioned clusters. The sample is an `/etc/haproxy/haproxy.cfg` configuration for an HAProxy load balancer. The example is not meant to provide advice for choosing one load balancing solution over another.

In the example, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

**NOTE**

If you are using HAProxy as a load balancer and SELinux is set to **enforcing**, you must ensure that the HAProxy service can bind to the configured TCP port by running **setsebool -P haproxy_connect_any=1**.

Example 18.6. Sample API and application Ingress load balancer configuration

```
global
  log      127.0.0.1 local2
  pidfile  /var/run/haproxy.pid
  maxconn  4000
  daemon
defaults
  mode          http
  log           global
  option        dontlognull
  option http-server-close
  option        redispatch
  retries       3
  timeout http-request 10s
  timeout queue 1m
  timeout connect 10s
  timeout client 1m
  timeout server 1m
  timeout http-keep-alive 10s
  timeout check 10s
  maxconn      3000
listen api-server-6443 1
  bind *:6443
  mode tcp
  option httpchk GET /readyz HTTP/1.0
  option log-health-checks
  balance roundrobin
  server bootstrap bootstrap.ocp4.example.com:6443 verify none check check-ssl inter 10s fall 2
  rise 3 backup 2
  server master0 master0.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
```

```

fall 2 rise 3
server master1 master1.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
fall 2 rise 3
server master2 master2.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
fall 2 rise 3
listen machine-config-server-22623 3
bind *:22623
mode tcp
server bootstrap bootstrap.ocp4.example.com:22623 check inter 1s backup 4
server master0 master0.ocp4.example.com:22623 check inter 1s
server master1 master1.ocp4.example.com:22623 check inter 1s
server master2 master2.ocp4.example.com:22623 check inter 1s
listen ingress-router-443 5
bind *:443
mode tcp
balance source
server compute0 compute0.ocp4.example.com:443 check inter 1s
server compute1 compute1.ocp4.example.com:443 check inter 1s
listen ingress-router-80 6
bind *:80
mode tcp
balance source
server compute0 compute0.ocp4.example.com:80 check inter 1s
server compute1 compute1.ocp4.example.com:80 check inter 1s

```

- 1 Port **6443** handles the Kubernetes API traffic and points to the control plane machines.
- 2 4 The bootstrap entries must be in place before the OpenShift Container Platform cluster installation and they must be removed after the bootstrap process is complete.
- 3 Port **22623** handles the machine config server traffic and points to the control plane machines.
- 5 Port **443** handles the HTTPS traffic and points to the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.
- 6 Port **80** handles the HTTP traffic and points to the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.



NOTE

If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application Ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes.

TIP

If you are using HAProxy as a load balancer, you can check that the **haproxy** process is listening on ports **6443**, **22623**, **443**, and **80** by running **netstat -nltp** on the HAProxy node.

18.3.5. Preparing the user-provisioned infrastructure

Before you install OpenShift Container Platform on user-provisioned infrastructure, you must prepare the underlying infrastructure.

This section provides details about the high-level steps required to set up your cluster infrastructure in preparation for an OpenShift Container Platform installation. This includes configuring IP networking and network connectivity for your cluster nodes, preparing a web server for the Ignition files, enabling the required ports through your firewall, and setting up the required DNS and load balancing infrastructure.

After preparation, your cluster infrastructure must meet the requirements outlined in the *Requirements for a cluster with user-provisioned infrastructure* section.

Prerequisites

- You have reviewed the [OpenShift Container Platform 4.x Tested Integrations](#) page.
- You have reviewed the infrastructure requirements detailed in the *Requirements for a cluster with user-provisioned infrastructure* section.

Procedure

1. Set up static IP addresses.
2. Set up an HTTP or HTTPS server to provide Ignition files to the cluster nodes.
3. Ensure that your network infrastructure provides the required network connectivity between the cluster components. See the *Networking requirements for user-provisioned infrastructure* section for details about the requirements.
4. Configure your firewall to enable the ports required for the OpenShift Container Platform cluster components to communicate. See *Networking requirements for user-provisioned infrastructure* section for details about the ports that are required.



IMPORTANT

By default, port **1936** is accessible for an OpenShift Container Platform cluster, because each control plane node needs access to this port.

Avoid using the Ingress load balancer to expose this port, because doing so might result in the exposure of sensitive information, such as statistics and metrics, related to Ingress Controllers.

5. Setup the required DNS infrastructure for your cluster.
 - a. Configure DNS name resolution for the Kubernetes API, the application wildcard, the bootstrap machine, the control plane machines, and the compute machines.
 - b. Configure reverse DNS resolution for the Kubernetes API, the bootstrap machine, the control plane machines, and the compute machines.
See the *User-provisioned DNS requirements* section for more information about the OpenShift Container Platform DNS requirements.
6. Validate your DNS configuration.
 - a. From your installation node, run DNS lookups against the record names of the Kubernetes API, the wildcard routes, and the cluster nodes. Validate that the IP addresses in the

responses correspond to the correct components.

- b. From your installation node, run reverse DNS lookups against the IP addresses of the load balancer and the cluster nodes. Validate that the record names in the responses correspond to the correct components.

See the *Validating DNS resolution for user-provisioned infrastructure* section for detailed DNS validation steps.

7. Provision the required API and application ingress load balancing infrastructure. See the *Load balancing requirements for user-provisioned infrastructure* section for more information about the requirements.



NOTE

Some load balancing solutions require the DNS name resolution for the cluster nodes to be in place before the load balancing is initialized.

18.3.6. Validating DNS resolution for user-provisioned infrastructure

You can validate your DNS configuration before installing OpenShift Container Platform on user-provisioned infrastructure.



IMPORTANT

The validation steps detailed in this section must succeed before you install your cluster.

Prerequisites

- You have configured the required DNS records for your user-provisioned infrastructure.

Procedure

1. From your installation node, run DNS lookups against the record names of the Kubernetes API, the wildcard routes, and the cluster nodes. Validate that the IP addresses contained in the responses correspond to the correct components.
 - a. Perform a lookup against the Kubernetes API record name. Check that the result points to the IP address of the API load balancer:

```
$ dig +noall +answer @<nameserver_ip> api.<cluster_name>.<base_domain> 1
```

- 1 Replace **<nameserver_ip>** with the IP address of the nameserver, **<cluster_name>** with your cluster name, and **<base_domain>** with your base domain name.

Example output

```
api.ocp4.example.com. 604800 IN A 192.168.1.5
```

- b. Perform a lookup against the Kubernetes internal API record name. Check that the result points to the IP address of the API load balancer:

```
$ dig +noall +answer @<nameserver_ip> api-int.<cluster_name>.<base_domain>
```

Example output

```
api-int.ocp4.example.com. 604800 IN A 192.168.1.5
```

- c. Test an example `*.apps.<cluster_name>.<base_domain>` DNS wildcard lookup. All of the application wildcard lookups must resolve to the IP address of the application ingress load balancer:

```
$ dig +noall +answer @<nameserver_ip> random.apps.<cluster_name>.<base_domain>
```

Example output

```
random.apps.ocp4.example.com. 604800 IN A 192.168.1.5
```



NOTE

In the example outputs, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

You can replace **random** with another wildcard value. For example, you can query the route to the OpenShift Container Platform console:

```
$ dig +noall +answer @<nameserver_ip> console-openshift-console.apps.<cluster_name>.<base_domain>
```

Example output

```
console-openshift-console.apps.ocp4.example.com. 604800 IN A 192.168.1.5
```

- d. Run a lookup against the bootstrap DNS record name. Check that the result points to the IP address of the bootstrap node:

```
$ dig +noall +answer @<nameserver_ip> bootstrap.<cluster_name>.<base_domain>
```

Example output

```
bootstrap.ocp4.example.com. 604800 IN A 192.168.1.96
```

- e. Use this method to perform lookups against the DNS record names for the control plane and compute nodes. Check that the results correspond to the IP addresses of each node.
2. From your installation node, run reverse DNS lookups against the IP addresses of the load balancer and the cluster nodes. Validate that the record names contained in the responses correspond to the correct components.
 - a. Perform a reverse lookup against the IP address of the API load balancer. Check that the response includes the record names for the Kubernetes API and the Kubernetes internal API:

```
$ dig +noall +answer @<nameserver_ip> -x 192.168.1.5
```

Example output

```
5.1.168.192.in-addr.arpa. 604800 IN PTR api-int.ocp4.example.com. 1
5.1.168.192.in-addr.arpa. 604800 IN PTR api.ocp4.example.com. 2
```

- 1 Provides the record name for the Kubernetes internal API.
- 2 Provides the record name for the Kubernetes API.



NOTE

A PTR record is not required for the OpenShift Container Platform application wildcard. No validation step is needed for reverse DNS resolution against the IP address of the application ingress load balancer.

- b. Perform a reverse lookup against the IP address of the bootstrap node. Check that the result points to the DNS record name of the bootstrap node:

```
$ dig +noall +answer @<nameserver_ip> -x 192.168.1.96
```

Example output

```
96.1.168.192.in-addr.arpa. 604800 IN PTR bootstrap.ocp4.example.com.
```

- c. Use this method to perform reverse lookups against the IP addresses for the control plane and compute nodes. Check that the results correspond to the DNS record names of each node.

18.3.7. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.



NOTE

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.



NOTE

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```



NOTE

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1 Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

18.3.8. Manually creating the installation configuration file

Installing the cluster requires that you manually create the installation configuration file.

Prerequisites

- You have an SSH public key on your local machine to provide to the installation program. The key will be used for SSH authentication onto your cluster nodes for debugging and disaster recovery.
- You have obtained the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create an installation directory to store your required installation assets in:

```
$ mkdir <installation_directory>
```



IMPORTANT

You must create a directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.


2. Customize the sample **install-config.yaml** file template that is provided and save it in the **<installation_directory>**.



NOTE

You must name this configuration file **install-config.yaml**.

3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the next step of the installation process. You must back it up now.

Additional resources

- [Installation configuration parameters for IBM Z®](#)

18.3.8.1. Sample install-config.yaml file for IBM Z

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.

```

apiVersion: v1
baseDomain: example.com 1
compute: 2
- hyperthreading: Enabled 3
  name: worker
  replicas: 0 4
  architecture: s390x
controlPlane: 5
  hyperthreading: Enabled 6
  name: master
  replicas: 3 7
  architecture: s390x
metadata:
  name: test 8
networking:
  clusterNetwork:
  - cidr: 10.128.0.0/14 9
    hostPrefix: 23 10
  networkType: OVNKubernetes 11
  serviceNetwork: 12
  - 172.30.0.0/16
platform:
  none: {} 13
fips: false 14
pullSecret: '{"auths":{"<local_registry>":{"auth": "<credentials>","email": "you@example.com"}}}' 15
sshKey: 'ssh-ed25519 AAAA...' 16
additionalTrustBundle: | 17
  -----BEGIN CERTIFICATE-----
  /-----END CERTIFICATE-----
imageContentSources: 18
- mirrors:
  - <local_repository>/ocp4/openshift4
  source: quay.io/openshift-release-dev/ocp-release
- mirrors:
  - <local_repository>/ocp4/openshift4
  source: quay.io/openshift-release-dev/ocp-v4.0-art-dev

```

1 The base domain of the cluster. All DNS records must be sub-domains of this base and include the


cluster name.

- 2 5 The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.
- 3 6 Specifies whether to enable or disable simultaneous multithreading (SMT), or hyperthreading. By default, SMT is enabled to increase the performance of the cores in your machines. You can disable it by setting the parameter value to **Disabled**. If you disable SMT, you must disable it in all cluster machines; this includes both control plane and compute machines.



NOTE

Simultaneous multithreading (SMT) is enabled by default. If SMT is not available on your OpenShift Container Platform nodes, the **hyperthreading** parameter has no effect.



IMPORTANT

If you disable **hyperthreading**, whether on your OpenShift Container Platform nodes or in the **install-config.yaml** file, ensure that your capacity planning accounts for the dramatically decreased machine performance.

- 4 You must set this value to **0** when you install OpenShift Container Platform on user-provisioned infrastructure. In installer-provisioned installations, the parameter controls the number of compute machines that the cluster creates and manages for you. In user-provisioned installations, you must manually deploy the compute machines before you finish installing the cluster.



NOTE

If you are installing a three-node cluster, do not deploy any compute machines when you install the Red Hat Enterprise Linux CoreOS (RHCOS) machines.

- 7 The number of control plane machines that you add to the cluster. Because the cluster uses these values as the number of etcd endpoints in the cluster, the value must match the number of control plane machines that you deploy.
- 8 The cluster name that you specified in your DNS records.
- 9 A block of IP addresses from which pod IP addresses are allocated. This block must not overlap with existing physical networks. These IP addresses are used for the pod network. If you need to access the pods from an external network, you must configure load balancers and routers to manage the traffic.



NOTE

Class E CIDR range is reserved for a future use. To use the Class E CIDR range, you must ensure your networking environment accepts the IP addresses within the Class E CIDR range.

- 10 The subnet prefix length to assign to each individual node. For example, if **hostPrefix** is set to **23**, then each node is assigned a **/23** subnet out of the given **cidr**, which allows for 510 ($2^{(32 - 23)} - 2$) pod IP addresses. If you are required to provide access to nodes from an external network,

configure load balancers and routers to manage the traffic.

- 11 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 12 The IP address pool to use for service IP addresses. You can enter only one IP address pool. This block must not overlap with existing physical networks. If you need to access the services from an external network, configure load balancers and routers to manage the traffic.
- 13 You must set the platform to **none**. You cannot provide additional platform configuration variables for IBM Z[®] infrastructure.



IMPORTANT

Clusters that are installed with the platform type **none** are unable to use some features, such as managing compute machines with the Machine API. This limitation applies even if the compute machines that are attached to the cluster are installed on a platform that would normally support the feature. This parameter cannot be changed after installation.

- 14 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.



IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 15 For **<local_registry>**, specify the registry domain name, and optionally the port, that your mirror registry uses to serve content. For example, **registry.example.com** or **registry.example.com:5000**. For **<credentials>**, specify the base64-encoded user name and password for your mirror registry.
- 16 The SSH public key for the **core** user in Red Hat Enterprise Linux CoreOS (RHCOS).



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- 17 Add the **additionalTrustBundle** parameter and value. The value must be the contents of the certificate file that you used for your mirror registry. The certificate file can be an existing, trusted certificate authority or the self-signed certificate that you generated for the mirror registry.

- 18 Provide the **imageContentSources** section according to the output of the command that you used to mirror the repository.



IMPORTANT

- When using the **oc adm release mirror** command, use the output from the **imageContentSources** section.
- When using **oc mirror** command, use the **repositoryDigestMirrors** section of the **ImageContentSourcePolicy** file that results from running the command.
- **ImageContentSourcePolicy** is deprecated. For more information see *Configuring image registry repository mirroring*.

18.3.8.2. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
```


```
<MY_TRUSTED_CA_CERT>
-----END CERTIFICATE-----
additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.



NOTE

The installation program does not support the proxy **readinessEndpoints** field.



NOTE

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.



NOTE

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

18.3.8.3. Configuring a three-node cluster

Optionally, you can deploy zero compute machines in a minimal three node cluster that consists of three control plane machines only. This provides smaller, more resource efficient clusters for cluster administrators and developers to use for testing, development, and production.

In three-node OpenShift Container Platform environments, the three control plane machines are schedulable, which means that your application workloads are scheduled to run on them.

Prerequisites

- You have an existing **install-config.yaml** file.

Procedure

- Ensure that the number of compute replicas is set to **0** in your **install-config.yaml** file, as shown in the following **compute** stanza:

```
compute:
- name: worker
  platform: {}
  replicas: 0
```



NOTE

You must set the value of the **replicas** parameter for the compute machines to **0** when you install OpenShift Container Platform on user-provisioned infrastructure, regardless of the number of compute machines you are deploying. In installer-provisioned installations, the parameter controls the number of compute machines that the cluster creates and manages for you. This does not apply to user-provisioned installations, where the compute machines are deployed manually.



NOTE

The preferred resource for control plane nodes is six vCPUs and 21 GB. For three control plane nodes this is the memory + vCPU equivalent of a minimum five-node cluster. You should back the three nodes, each installed on a 120 GB disk, with three IFLs that are SMT2 enabled. The minimum tested setup is three vCPUs and 10 GB on a 120 GB disk for each control plane node.

For three-node cluster installations, follow these next steps:

- If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes. See the *Load balancing requirements for user-provisioned infrastructure* section for more information.
- When you create the Kubernetes manifest files in the following procedure, ensure that the **mastersSchedulable** parameter in the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** file is set to **true**. This enables your application workloads to run on the control plane nodes.
- Do not deploy any compute nodes when you create the Red Hat Enterprise Linux CoreOS (RHCOS) machines.

18.3.9. Cluster Network Operator configuration

The configuration for the cluster network is specified as part of the Cluster Network Operator (CNO) configuration and stored in a custom resource (CR) object that is named **cluster**. The CR specifies the fields for the **Network** API in the **operator.openshift.io** API group.

The CNO configuration inherits the following fields during cluster installation from the **Network** API in the **Network.config.openshift.io** API group:

clusterNetwork

IP address pools from which pod IP addresses are allocated.

serviceNetwork

IP address pool for services.

defaultNetwork.type

Cluster network plugin. **OVNKubernetes** is the only supported plugin during installation.

You can specify the cluster network plugin configuration for your cluster by setting the fields for the **defaultNetwork** object in the CNO object named **cluster**.

18.3.9.1. Cluster Network Operator configuration object

The fields for the Cluster Network Operator (CNO) are described in the following table:

Table 18.29. Cluster Network Operator configuration object


Field	Type	Description
metadata.name	string	The name of the CNO object. This name is always cluster .
spec.clusterNetwork	array	A list specifying the blocks of IP addresses from which pod IP addresses are allocated and the subnet prefix length assigned to each individual node in the cluster. For example: <pre>spec: clusterNetwork: - cidr: 10.128.0.0/19 hostPrefix: 23 - cidr: 10.128.32.0/19 hostPrefix: 23</pre>

Field	Type	Description
spec.serviceNetwork	array	<p>A block of IP addresses for services. The OpenShift SDN and OVN-Kubernetes network plugins support only a single IP address block for the service network. For example:</p> <pre>spec: serviceNetwork: - 172.30.0.0/14</pre> <p>You can customize this field only in the install-config.yaml file before you create the manifests. The value is read-only in the manifest file.</p>
spec.defaultNetwork	object	Configures the network plugin for the cluster network.
spec.kubeProxyConfig	object	The fields for this object specify the kube-proxy configuration. If you are using the OVN-Kubernetes cluster network plugin, the kube-proxy configuration has no effect.

defaultNetwork object configuration

The values for the **defaultNetwork** object are defined in the following table:

Table 18.30. defaultNetwork object

Field	Type	Description
type	string	<p>OVNKubernetes. The Red Hat OpenShift Networking network plugin is selected during installation. This value cannot be changed after cluster installation.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p>NOTE</p> <p>OpenShift Container Platform uses the OVN-Kubernetes network plugin by default. OpenShift SDN is no longer available as an installation choice for new clusters.</p> </div> </div>
ovnKubernetesConfig	object	This object is only valid for the OVN-Kubernetes network plugin.

Configuration for the OVN-Kubernetes network plugin

The following table describes the configuration fields for the OVN-Kubernetes network plugin:

Table 18.31. ovnKubernetesConfig object


Field	Type	Description
mtu	integer	<p>The maximum transmission unit (MTU) for the Geneve (Generic Network Virtualization Encapsulation) overlay network. This is detected automatically based on the MTU of the primary network interface. You do not normally need to override the detected MTU.</p> <p>If the auto-detected value is not what you expect it to be, confirm that the MTU on the primary network interface on your nodes is correct. You cannot use this option to change the MTU value of the primary network interface on the nodes.</p> <p>If your cluster requires different MTU values for different nodes, you must set this value to 100 less than the lowest MTU value in your cluster. For example, if some nodes in your cluster have an MTU of 9001, and some have an MTU of 1500, you must set this value to 1400.</p>
genevePort	integer	The port to use for all Geneve packets. The default value is 6081 . This value cannot be changed after cluster installation.
ipsecConfig	object	Specify a configuration object for customizing the IPsec configuration.
ipv4	object	Specifies a configuration object for IPv4 settings.
ipv6	object	Specifies a configuration object for IPv6 settings.
policyAuditConfig	object	Specify a configuration object for customizing network policy audit logging. If unset, the defaults audit log settings are used.
gatewayConfig	object	<p>Optional: Specify a configuration object for customizing how egress traffic is sent to the node gateway.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p>NOTE</p> <p>While migrating egress traffic, you can expect some disruption to workloads and service traffic until the Cluster Network Operator (CNO) successfully rolls out the changes.</p> </div> </div>

Table 18.32. `ovnKubernetesConfig.ipv4` object

Field	Type	Description
-------	------	-------------

Field	Type	Description
internalTransitSwitchSubnet	string	<p>If your existing network infrastructure overlaps with the 100.88.0.0/16 IPv4 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. The subnet for the distributed transit switch that enables east-west traffic. This subnet cannot overlap with any other subnets used by OVN-Kubernetes or on the host itself. It must be large enough to accommodate one IP address per node in your cluster.</p> <p>The default value is 100.88.0.0/16.</p>
internalJoinSubnet	string	<p>If your existing network infrastructure overlaps with the 100.64.0.0/16 IPv4 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. You must ensure that the IP address range does not overlap with any other subnet used by your OpenShift Container Platform installation. The IP address range must be larger than the maximum number of nodes that can be added to the cluster. For example, if the clusterNetwork.cidr value is 10.128.0.0/14 and the clusterNetwork.hostPrefix value is /23, then the maximum number of nodes is 2⁽²³⁻¹⁴⁾=512.</p> <p>The default value is 100.64.0.0/16.</p>

Table 18.33. `ovnKubernetesConfig.ipv6` object

Field	Type	Description
internalTransitSwitchSubnet	string	<p>If your existing network infrastructure overlaps with the fd97::<64 IPv6 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. The subnet for the distributed transit switch that enables east-west traffic. This subnet cannot overlap with any other subnets used by OVN-Kubernetes or on the host itself. It must be large enough to accommodate one IP address per node in your cluster.</p> <p>The default value is fd97::<64.</p>
internalJoinSubnet	string	<p>If your existing network infrastructure overlaps with the fd98::<64 IPv6 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. You must ensure that the IP address range does not overlap with any other subnet used by your OpenShift Container Platform installation. The IP address range must be larger than the maximum number of nodes that can be added to the cluster.</p> <p>The default value is fd98::<64.</p>

Table 18.34. `policyAuditConfig` object

Field	Type	Description
rateLimit	integer	The maximum number of messages to generate every second per node. The default value is 20 messages per second.
maxFileSize	integer	The maximum size for the audit log in bytes. The default value is 50000000 or 50 MB.
maxLogFiles	integer	The maximum number of log files that are retained.
destination	string	One of the following additional audit log targets: libc The libc syslog() function of the journald process on the host. udp:<host>:<port> A syslog server. Replace <host>:<port> with the host and port of the syslog server. unix:<file> A Unix Domain Socket file specified by <file> . null Do not send the audit logs to any additional target.
syslogFacility	string	The syslog facility, such as kern , as defined by RFC5424. The default value is local0 .

Table 18.35. gatewayConfig object

Field	Type	Description
routingViaHost	boolean	Set this field to true to send egress traffic from pods to the host networking stack. For highly-specialized installations and applications that rely on manually configured routes in the kernel routing table, you might want to route egress traffic to the host networking stack. By default, egress traffic is processed in OVN to exit the cluster and is not affected by specialized routes in the kernel routing table. The default value is false . This field has an interaction with the Open vSwitch hardware offloading feature. If you set this field to true , you do not receive the performance benefits of the offloading because egress traffic is processed by the host networking stack.
ipForwarding	object	You can control IP forwarding for all traffic on OVN-Kubernetes managed interfaces by using the ipForwarding specification in the Network resource. Specify Restricted to only allow IP forwarding for Kubernetes related traffic. Specify Global to allow forwarding of all IP traffic. For new installations, the default is Restricted . For updates to OpenShift Container Platform 4.14 or later, the default is Global .

Field	Type	Description
ipv4	object	Optional: Specify an object to configure the internal OVN-Kubernetes masquerade address for host to service traffic for IPv4 addresses.
ipv6	object	Optional: Specify an object to configure the internal OVN-Kubernetes masquerade address for host to service traffic for IPv6 addresses.

Table 18.36. gatewayConfig.ipv4 object

Field	Type	Description
internalMasqueradeSubnet	string	The masquerade IPv4 addresses that are used internally to enable host to service traffic. The host is configured with these IP addresses as well as the shared gateway bridge interface. The default value is 169.254.169.0/29 .

Table 18.37. gatewayConfig.ipv6 object

Field	Type	Description
internalMasqueradeSubnet	string	The masquerade IPv6 addresses that are used internally to enable host to service traffic. The host is configured with these IP addresses as well as the shared gateway bridge interface. The default value is fd69::/125 .

Table 18.38. ipsecConfig object

Field	Type	Description
mode	string	Specifies the behavior of the IPsec implementation. Must be one of the following values: <ul style="list-style-type: none"> ● Disabled: IPsec is not enabled on cluster nodes. ● External: IPsec is enabled for network traffic with external hosts. ● Full: IPsec is enabled for pod traffic and network traffic with external hosts.

Example OVN-Kubernetes configuration with IPsec enabled

```
defaultNetwork:
  type: OVNKubernetes
  ovnKubernetesConfig:
    mtu: 1400
```

```
genevePort: 6081
ipsecConfig:
mode: Full
```


**IMPORTANT**

Using OVNKubernetes can lead to a stack exhaustion problem on IBM Power®.

kubeProxyConfig object configuration (OpenShiftSDN container network interface only)

The values for the **kubeProxyConfig** object are defined in the following table:

Table 18.39. kubeProxyConfig object

Field	Type	Description
iptablesSyncPeriod	string	<p>The refresh period for iptables rules. The default value is 30s. Valid suffixes include s, m, and h and are described in the Go time package documentation.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p>NOTE</p> <p>Because of performance improvements introduced in OpenShift Container Platform 4.3 and greater, adjusting the iptablesSyncPeriod parameter is no longer necessary.</p> </div> </div>
proxyArguments.iptables-min-sync-period	array	<p>The minimum duration before refreshing iptables rules. This field ensures that the refresh does not happen too frequently. Valid suffixes include s, m, and h and are described in the Go time package. The default value is:</p> <pre>kubeProxyConfig: proxyArguments: iptables-min-sync-period: - 0s</pre>

18.3.10. Creating the Kubernetes manifest and Ignition config files

Because you must modify some cluster definition files and manually start the cluster machines, you must generate the Kubernetes manifest and Ignition config files that the cluster needs to configure the machines.

The installation configuration file transforms into the Kubernetes manifests. The manifests wrap into the Ignition configuration files, which are later used to configure the cluster machines.



IMPORTANT

- The Ignition config files that the OpenShift Container Platform installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.



NOTE

The installation program that generates the manifest and Ignition files is architecture specific and can be obtained from the [client image mirror](#). The Linux version of the installation program runs on s390x only. This installer program is also available as a Mac OS version.

Prerequisites

- You obtained the OpenShift Container Platform installation program. For a restricted network installation, these files are on your mirror host.
- You created the **install-config.yaml** installation configuration file.

Procedure

1. Change to the directory that contains the OpenShift Container Platform installation program and generate the Kubernetes manifests for the cluster:


```
$ ./openshift-install create manifests --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the installation directory that contains the **install-config.yaml** file you created.



WARNING

If you are installing a three-node cluster, skip the following step to allow the control plane nodes to be schedulable.



IMPORTANT

When you configure control plane nodes from the default unschedulable to schedulable, additional subscriptions are required. This is because control plane nodes then become compute nodes.

2. Check that the **mastersSchedulable** parameter in the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** Kubernetes manifest file is set to **false**. This setting prevents pods from being scheduled on the control plane machines:
 - a. Open the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** file.
 - b. Locate the **mastersSchedulable** parameter and ensure that it is set to **false**.
 - c. Save and exit the file.
3. To create the Ignition configuration files, run the following command from the directory that contains the installation program:

```
$ ./openshift-install create ignition-configs --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the same installation directory.

Ignition config files are created for the bootstrap, control plane, and compute nodes in the installation directory. The **kubeadmin-password** and **kubeconfig** files are created in the **./<installation_directory>/auth** directory:

```
.
├── auth
│   ├── kubeadmin-password
│   └── kubeconfig
├── bootstrap.ign
├── master.ign
├── metadata.json
└── worker.ign
```

18.3.11. Configuring NBDE with static IP in an IBM Z or IBM LinuxONE environment

Enabling NBDE disk encryption in an IBM Z® or IBM® LinuxONE environment requires additional steps, which are described in detail in this section.

Prerequisites

- You have set up the External Tang Server. See [Network-bound disk encryption](#) for instructions.
- You have installed the **butane** utility.
- You have reviewed the instructions for how to create machine configs with Butane.

Procedure

1. Create Butane configuration files for the control plane and compute nodes.

The following example of a Butane configuration for a control plane node creates a file named **master-storage.bu** for disk encryption:

```
variant: openshift
version: 4.16.0
metadata:
  name: master-storage
  labels:
    machineconfiguration.openshift.io/role: master
storage:
  luks:
    - clevis:
      tang:
        - thumbprint: QcPr_NHFJammnRCA3fFMVdNBwjs
          url: http://clevis.example.com:7500
      options: ❶
        - --cipher
        - aes-cbc-essiv:sha256
      device: /dev/disk/by-partlabel/root ❷
      label: luks-root
      name: root
      wipe_volume: true
    filesystems:
      - device: /dev/mapper/root
        format: xfs
        label: root
        wipe_filesystem: true
  openshift:
    fips: true ❸
```

- ❶ The cipher option is only required if FIPS mode is enabled. Omit the entry if FIPS is disabled.
- ❷ For installations on DASD-type disks, replace with **device: /dev/disk/by-label/root**.
- ❸ Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.

2. Create a customized initramfs file to boot the machine, by running the following command:

```
$ coreos-installer pxe customize \
  /root/rhcos-bootfiles/rhcos-<release>-live-initramfs.s390x.img \
  --dest-device /dev/disk/by-id/scsi-<serial_number> --dest-karg-append \
  ip=<ip_address>::<gateway_ip>:<subnet_mask>::<network_device>:none \
  --dest-karg-append nameserver=<nameserver_ip> \
  --dest-karg-append rd.neednet=1 -o \
  /root/rhcos-bootfiles/<node_name>-initramfs.s390x.img
```



NOTE

Before first boot, you must customize the initramfs for each node in the cluster, and add PXE kernel parameters.

3. Create a parameter file that includes **ignition.platform.id=metal** and **ignition.firstboot**.

Example kernel parameter file for the control plane machine:

```
rd.neednet=1 \
console=ttySCLP0 \
coreos.inst.install_dev=/dev/dasda \ 1
ignition.firstboot ignition.platform.id=metal \
coreos.live.rootfs_url=http://<http_server>/rhcos-<version>-live-rootfs.<architecture>.img \ 2
coreos.inst.ignition_url=http://<http_server>/master.ign \ 3
ip=10.19.17.2::10.19.17.1:255.255.255.0::enb0d0:none nameserver=10.19.17.1 \
zfcplib.allow_lun_scan=0 \ 4
rd.znet=qeth,0.0.b0d0,0.0.b0d1,0.0.b0d2,layer2=1 \
rd.zfcplib=0.0.5677,0x600606680g7f0056,0x034F000000000000 \ 5
```

- 1 For installations on DASD-type disks, add **coreos.inst.install_dev=/dev/dasda**. Omit this value for FCP-type disks.
- 2 Specify the location of the **rootfs** artifact for the **kernel** and **initramfs** you are booting. Only HTTP and HTTPS protocols are supported.
- 3 Specify the location of the Ignition config file. Use **master.ign** or **worker.ign**. Only HTTP and HTTPS protocols are supported.
- 4 For installations on FCP-type disks, add **zfcplib.allow_lun_scan=0**. Omit this value for DASD-type disks.
- 5 For installations on DASD-type disks, replace with **rd.dasd=0.0.3490** to specify the DASD device.



NOTE

Write all options in the parameter file as a single line and make sure you have no newline characters.

Additional resources

- [Creating machine configs with Butane](#)

18.3.12. Installing RHCOS and starting the OpenShift Container Platform bootstrap process

To install OpenShift Container Platform on IBM Z® infrastructure that you provision, you must install Red Hat Enterprise Linux CoreOS (RHCOS) on z/VM guest virtual machines. When you install RHCOS, you must provide the Ignition config file that was generated by the OpenShift Container Platform installation program for the type of machine you are installing. If you have configured suitable networking, DNS, and load balancing infrastructure, the OpenShift Container Platform bootstrap process begins automatically after the RHCOS z/VM guest virtual machines have rebooted.

Complete the following steps to create the machines.

Prerequisites

- An HTTP or HTTPS server running on your provisioning machine that is accessible to the machines you create.
- If you want to enable secure boot, you have obtained the appropriate Red Hat Product Signing Key and read [Secure boot on IBM Z and IBM LinuxONE](#) in IBM documentation.

Procedure

1. Log in to Linux on your provisioning machine.
2. Obtain the Red Hat Enterprise Linux CoreOS (RHCOS) kernel, initramfs, and rootfs files from the [RHCOS image mirror](#).



IMPORTANT

The RHCOS images might not change with every release of OpenShift Container Platform. You must download images with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Only use the appropriate kernel, initramfs, and rootfs artifacts described in the following procedure.

The file names contain the OpenShift Container Platform version number. They resemble the following examples:

- kernel: **rhcos-`<version>`-live-kernel-`<architecture>`**
- initramfs: **rhcos-`<version>`-live-initramfs.`<architecture>`.img**
- rootfs: **rhcos-`<version>`-live-rootfs.`<architecture>`.img**



NOTE

The rootfs image is the same for FCP and DASD.

3. Create parameter files. The following parameters are specific for a particular virtual machine:
 - For **ip=**, specify the following seven entries:
 - i. The IP address for the machine.
 - ii. An empty string.
 - iii. The gateway.
 - iv. The netmask.
 - v. The machine host and domain name in the form **hostname.domainname**. Omit this value to let RHCOS decide.
 - vi. The network interface name. Omit this value to let RHCOS decide.
 - vii. If you use static IP addresses, specify **none**.
 - For **coreos.inst.ignition_url=**, specify the Ignition file for the machine role. Use **bootstrap.ign**, **master.ign**, or **worker.ign**. Only HTTP and HTTPS protocols are supported.

- For **coreos.live.rootfs_url=**, specify the matching rootfs artifact for the kernel and initramfs you are booting. Only HTTP and HTTPS protocols are supported.
- Optional: To enable secure boot, add **coreos.inst.secure_ign**
- For installations on DASD-type disks, complete the following tasks:
 - i. For **coreos.inst.install_dev=**, specify **/dev/dasda**.
 - ii. Use **rd.dasd=** to specify the DASD where RHCOS is to be installed.
 - iii. Leave all other parameters unchanged.

Example parameter file, **bootstrap-0.parm**, for the bootstrap machine:

```
rd.neednet=1 \
console=ttysclp0 \
coreos.inst.install_dev=/dev/dasda \
coreos.live.rootfs_url=http://<http_server>/rhcos-<version>-live-rootfs.
<architecture>.img \ 1
coreos.inst.ignition_url=http://<http_server>/bootstrap.ign \ 2
coreos.inst.secure_ign \ 3
ip=172.18.78.2::172.18.78.1:255.255.255.0::none nameserver=172.18.78.1 \
rd.znet=qeth,0.0.bdf0,0.0.bdf1,0.0.bdf2,layer2=1,portno=0 \
zfcplib.allow_lun_scan=0 \
rd.dasd=0.0.3490
```

- 1 Specify the location of the **rootfs** artifact for the **kernel** and **initramfs** you are booting. Only HTTP and HTTPS protocols are supported.
- 2 Specify the location of the Ignition config file. Use **bootstrap.ign**, **master.ign**, or **worker.ign**. Only HTTP and HTTPS protocols are supported.
- 3 Optional: To enable secure boot, add **coreos.inst.secure_ign**.

Write all options in the parameter file as a single line and make sure you have no newline characters.

- For installations on FCP-type disks, complete the following tasks:
 - i. Use **rd.zfcplib=<adapter>,<wwpn>,<lun>** to specify the FCP disk where RHCOS is to be installed. For multipathing repeat this step for each additional path.



NOTE

When you install with multiple paths, you must enable multipathing directly after the installation, not at a later point in time, as this can cause problems.

- ii. Set the install device as: **coreos.inst.install_dev=/dev/disk/by-id/scsi-<serial_number>**.

**NOTE**

If additional LUNs are configured with NPIV, FCP requires **zfcplib.allow_lun_scan=0**. If you must enable **zfcplib.allow_lun_scan=1** because you use a CSI driver, for example, you must configure your NPIV so that each node cannot access the boot partition of another node.

- iii. Leave all other parameters unchanged.

**IMPORTANT**

Additional postinstallation steps are required to fully enable multipathing. For more information, see "Enabling multipathing with kernel arguments on RHCOS" in *Postinstallation machine configuration tasks*.

The following is an example parameter file **worker-1.parm** for a compute node with multipathing:

```
rd.neednet=1 \
console=ttysclp0 \
coreos.inst.install_dev=/dev/disk/by-id/scsi-<serial_number> \
coreos.live.rootfs_url=http://<http_server>/rhcos-<version>-live-rootfs.
<architecture>.img \
coreos.inst.ignition_url=http://<http_server>/worker.ign \
ip=172.18.78.2::172.18.78.1:255.255.255.0::none nameserver=172.18.78.1 \
rd.znet=qeth,0.0.bdf0,0.0.bdf1,0.0.bdf2,layer2=1,portno=0 \
zfcplib.allow_lun_scan=0 \
rd.zfcplib=0.0.1987,0x50050763070bc5e3,0x4008400B00000000 \
rd.zfcplib=0.0.19C7,0x50050763070bc5e3,0x4008400B00000000 \
rd.zfcplib=0.0.1987,0x50050763071bc5e3,0x4008400B00000000 \
rd.zfcplib=0.0.19C7,0x50050763071bc5e3,0x4008400B00000000
```

Write all options in the parameter file as a single line and make sure you have no newline characters.

4. Transfer the initramfs, kernel, parameter files, and RHCOS images to z/VM, for example with FTP. For details about how to transfer the files with FTP and boot from the virtual reader, see [Installing under Z/VM](#).
5. Punch the files to the virtual reader of the z/VM guest virtual machine that is to become your bootstrap node.
See [PUNCH](#) in IBM Documentation.

TIP

You can use the CP PUNCH command or, if you use Linux, the **vmur** command to transfer files between two z/VM guest virtual machines.

6. Log in to CMS on the bootstrap machine.
7. IPL the bootstrap machine from the reader:

```
$ ipl c
```

See [IPL](#) in IBM Documentation.


8. Repeat this procedure for the other machines in the cluster.

18.3.12.1. Advanced RHCOS installation reference

This section illustrates the networking configuration and other advanced options that allow you to modify the Red Hat Enterprise Linux CoreOS (RHCOS) manual installation process. The following tables describe the kernel arguments and command-line options you can use with the RHCOS live installer and the **coreos-installer** command.

18.3.12.1.1. Networking and bonding options for ISO installations

If you install RHCOS from an ISO image, you can add kernel arguments manually when you boot the image to configure networking for a node. If no networking arguments are specified, DHCP is activated in the initramfs when RHCOS detects that networking is required to fetch the Ignition config file.



IMPORTANT

When adding networking arguments manually, you must also add the **rd.neednet=1** kernel argument to bring the network up in the initramfs.

The following information provides examples for configuring networking and bonding on your RHCOS nodes for ISO installations. The examples describe how to use the **ip=**, **nameserver=**, and **bond=** kernel arguments.



NOTE

Ordering is important when adding the kernel arguments: **ip=**, **nameserver=**, and then **bond=**.

The networking options are passed to the **dracut** tool during system boot. For more information about the networking options supported by **dracut**, see the [dracut.cmdline manual page](#).

The following examples are the networking options for ISO installation.

Configuring DHCP or static IP addresses

To configure an IP address, either use DHCP (**ip=dhcp**) or set an individual static IP address (**ip=<host_ip>**). If setting a static IP, you must then identify the DNS server IP address (**nameserver=<dns_ip>**) on each node. The following example sets:

- The node's IP address to **10.10.10.2**
- The gateway address to **10.10.10.254**
- The netmask to **255.255.255.0**
- The hostname to **core0.example.com**
- The DNS server address to **4.4.4.41**
- The auto-configuration value to **none**. No auto-configuration is required when IP networking is configured statically.

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none
nameserver=4.4.4.41
```



NOTE

When you use DHCP to configure IP addressing for the RHCOS machines, the machines also obtain the DNS server information through DHCP. For DHCP-based deployments, you can define the DNS server address that is used by the RHCOS nodes through your DHCP server configuration.

Configuring an IP address without a static hostname

You can configure an IP address without assigning a static hostname. If a static hostname is not set by the user, it will be picked up and automatically set by a reverse DNS lookup. To configure an IP address without a static hostname refer to the following example:

- The node's IP address to **10.10.10.2**
- The gateway address to **10.10.10.254**
- The netmask to **255.255.255.0**
- The DNS server address to **4.4.4.41**
- The auto-configuration value to **none**. No auto-configuration is required when IP networking is configured statically.

```
ip=10.10.10.2::10.10.10.254:255.255.255.0::enp1s0:none
nameserver=4.4.4.41
```

Specifying multiple network interfaces

You can specify multiple network interfaces by setting multiple **ip=** entries.

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none
ip=10.10.10.3::10.10.10.254:255.255.255.0:core0.example.com:enp2s0:none
```

Configuring default gateway and route

Optional: You can configure routes to additional networks by setting an **rd.route=** value.



NOTE

When you configure one or multiple networks, one default gateway is required. If the additional network gateway is different from the primary network gateway, the default gateway must be the primary network gateway.

- Run the following command to configure the default gateway:

```
ip=::10.10.10.254:::
```

- Enter the following command to configure the route for the additional network:

```
rd.route=20.20.20.0/24:20.20.20.254:enp2s0
```

Disabling DHCP on a single interface

You can disable DHCP on a single interface, such as when there are two or more network interfaces and only one interface is being used. In the example, the **enp1s0** interface has a static networking configuration and DHCP is disabled for **enp2s0**, which is not used:

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none
ip=::::core0.example.com:enp2s0:none
```

Combining DHCP and static IP configurations

You can combine DHCP and static IP configurations on systems with multiple network interfaces, for example:

```
ip=enp1s0:dhcp
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp2s0:none
```

Configuring VLANs on individual interfaces

Optional: You can configure VLANs on individual interfaces by using the **vlan=** parameter.

- To configure a VLAN on a network interface and use a static IP address, run the following command:

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp2s0.100:none
vlan=enp2s0.100:enp2s0
```

- To configure a VLAN on a network interface and to use DHCP, run the following command:

```
ip=enp2s0.100:dhcp
vlan=enp2s0.100:enp2s0
```

Providing multiple DNS servers

You can provide multiple DNS servers by adding a **nameserver=** entry for each server, for example:

```
nameserver=1.1.1.1
nameserver=8.8.8.8
```

Bonding multiple network interfaces to a single interface

Optional: You can bond multiple network interfaces to a single interface by using the **bond=** option. Refer to the following examples:

- The syntax for configuring a bonded interface is: **bond=<name>[:<network_interfaces>] [:options]**
<name> is the bonding device name (**bond0**), **<network_interfaces>** represents a comma-separated list of physical (ethernet) interfaces (**em1,em2**), and *options* is a comma-separated list of bonding options. Enter **modinfo bonding** to see available options.
- When you create a bonded interface using **bond=**, you must specify how the IP address is assigned and other information for the bonded interface.
 - To configure the bonded interface to use DHCP, set the bond's IP address to **dhcp**. For example:

```
bond=bond0:em1,em2:mode=active-backup
ip=bond0:dhcp
```

- To configure the bonded interface to use a static IP address, enter the specific IP address you want and related information. For example:

```
bond=bond0:em1,em2:mode=active-backup,fail_over_mac=1
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:bond0:none
```

Always set the **fail_over_mac=1** option in active-backup mode, to avoid problems when shared OSA/RoCE cards are used.

Bonding multiple network interfaces to a single interface

Optional: You can configure VLANs on bonded interfaces by using the **vlan=** parameter and to use DHCP, for example:

```
ip=bond0.100:dhcp
bond=bond0:em1,em2:mode=active-backup
vlan=bond0.100:bond0
```

Use the following example to configure the bonded interface with a VLAN and to use a static IP address:

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:bond0.100:none
bond=bond0:em1,em2:mode=active-backup
vlan=bond0.100:bond0
```

Using network teaming

Optional: You can use a network teaming as an alternative to bonding by using the **team=** parameter:

- The syntax for configuring a team interface is: **team=name[:network_interfaces]**
name is the team device name (**team0**) and *network_interfaces* represents a comma-separated list of physical (ethernet) interfaces (**em1, em2**).



NOTE

Teaming is planned to be deprecated when RHCOS switches to an upcoming version of RHEL. For more information, see this [Red Hat Knowledgebase Article](#).

Use the following example to configure a network team:

```
team=team0:em1,em2
ip=team0:dhcp
```

18.3.13. Waiting for the bootstrap process to complete

The OpenShift Container Platform bootstrap process begins after the cluster nodes first boot into the persistent RHCOS environment that has been installed to disk. The configuration information provided through the Ignition config files is used to initialize the bootstrap process and install OpenShift Container Platform on the machines. You must wait for the bootstrap process to complete.

Prerequisites

- You have created the Ignition config files for your cluster.
- You have configured suitable network, DNS and load balancing infrastructure.

- You have obtained the installation program and generated the Ignition config files for your cluster.
- You installed RHCOS on your cluster machines and provided the Ignition config files that the OpenShift Container Platform installation program generated.

Procedure

1. Monitor the bootstrap process:

```
$ ./openshift-install --dir <installation_directory> wait-for bootstrap-complete \ 1  
--log-level=info 2
```

1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Example output

```
INFO Waiting up to 30m0s for the Kubernetes API at https://api.test.example.com:6443...  
INFO API v1.29.4 up  
INFO Waiting up to 30m0s for bootstrapping to complete...  
INFO It is now safe to remove the bootstrap resources
```

The command succeeds when the Kubernetes API server signals that it has been bootstrapped on the control plane machines.

2. After the bootstrap process is complete, remove the bootstrap machine from the load balancer.



IMPORTANT

You must remove the bootstrap machine from the load balancer at this point. You can also remove or reformat the bootstrap machine itself.

18.3.14. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1 For `<installation_directory>`, specify the path to the directory that you stored the installation files in.

2. Verify you can run `oc` commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

18.3.15. Approving the certificate signing requests for your machines

When you add machines to a cluster, two pending certificate signing requests (CSRs) are generated for each machine that you added. You must confirm that these CSRs are approved or, if necessary, approve them yourself. The client requests must be approved first, followed by the server requests.

Prerequisites

- You added machines to your cluster.

Procedure

1. Confirm that the cluster recognizes the machines:

```
$ oc get nodes
```

Example output

```
NAME      STATUS   ROLES    AGE   VERSION
master-0  Ready    master   63m   v1.29.4
master-1  Ready    master   63m   v1.29.4
master-2  Ready    master   64m   v1.29.4
```

The output lists all of the machines that you created.



NOTE

The preceding output might not include the compute nodes, also known as worker nodes, until some CSRs are approved.

2. Review the pending CSRs and ensure that you see the client requests with the **Pending** or **Approved** status for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

```
NAME      AGE   REQUESTOR                                     CONDITION
csr-8b2br  15m   system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper  Pending
```

```
csr-8vnps 15m system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper Pending
...
```

In this example, two machines are joining the cluster. You might see more approved CSRs in the list.

- If the CSRs were not approved, after all of the pending CSRs for the machines you added are in **Pending** status, approve the CSRs for your cluster machines:



NOTE

Because the CSRs rotate automatically, approve your CSRs within an hour of adding the machines to the cluster. If you do not approve them within an hour, the certificates will rotate, and more than two certificates will be present for each node. You must approve all of these certificates. After the client CSR is approved, the Kubelet creates a secondary CSR for the serving certificate, which requires manual approval. Then, subsequent serving certificate renewal requests are automatically approved by the **machine-approver** if the Kubelet requests a new certificate with identical parameters.



NOTE

For clusters running on platforms that are not machine API enabled, such as bare metal and other user-provisioned infrastructure, you must implement a method of automatically approving the kubelet serving certificate requests (CSRs). If a request is not approved, then the **oc exec**, **oc rsh**, and **oc logs** commands cannot succeed, because a serving certificate is required when the API server connects to the kubelet. Any operation that contacts the Kubelet endpoint requires this certificate approval to be in place. The method must watch for new CSRs, confirm that the CSR was submitted by the **node-bootstrapper** service account in the **system:node** or **system:admin** groups, and confirm the identity of the node.

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{"\n"}\n{{end}}{{end}}' | xargs --no-run-if-empty oc adm certificate approve
```



NOTE

Some Operators might not become available until some CSRs are approved.

- Now that your client requests are approved, you must review the server requests for each machine that you added to the cluster:


```
$ oc get csr
```

Example output

```
NAME      AGE   REQUESTOR                                     CONDITION
csr-bfd72 5m26s system:node:ip-10-0-50-126.us-east-2.compute.internal
Pending
csr-c57lv 5m26s system:node:ip-10-0-95-157.us-east-2.compute.internal
Pending
...
```

- If the remaining CSRs are not approved, and are in the **Pending** status, approve the CSRs for your cluster machines:

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{"\n"}\n{{end}}{{end}}' | xargs oc adm certificate approve
```

- After all client and server CSRs have been approved, the machines have the **Ready** status. Verify this by running the following command:

```
$ oc get nodes
```

Example output

```
NAME      STATUS   ROLES    AGE   VERSION
master-0  Ready   master   73m   v1.29.4
master-1  Ready   master   73m   v1.29.4
master-2  Ready   master   74m   v1.29.4
worker-0  Ready   worker   11m   v1.29.4
worker-1  Ready   worker   11m   v1.29.4
```



NOTE

It can take a few minutes after approval of the server CSRs for the machines to transition to the **Ready** status.

Additional information

- For more information on CSRs, see [Certificate Signing Requests](#).

18.3.16. Initial Operator configuration

After the control plane initializes, you must immediately configure some Operators so that they all become available.

Prerequisites

- Your control plane has initialized.

Procedure

1. Watch the cluster components come online:

```
$ watch -n5 oc get clusteroperators
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.16.0	True	False	False	19m
baremetal	4.16.0	True	False	False	37m
cloud-credential	4.16.0	True	False	False	40m
cluster-autoscaler	4.16.0	True	False	False	37m
config-operator	4.16.0	True	False	False	38m
console	4.16.0	True	False	False	26m
csi-snapshot-controller	4.16.0	True	False	False	37m
dns	4.16.0	True	False	False	37m
etcd	4.16.0	True	False	False	36m
image-registry	4.16.0	True	False	False	31m
ingress	4.16.0	True	False	False	30m
insights	4.16.0	True	False	False	31m
kube-apiserver	4.16.0	True	False	False	26m
kube-controller-manager	4.16.0	True	False	False	36m
kube-scheduler	4.16.0	True	False	False	36m
kube-storage-version-migrator	4.16.0	True	False	False	37m
machine-api	4.16.0	True	False	False	29m
machine-approver	4.16.0	True	False	False	37m
machine-config	4.16.0	True	False	False	36m
marketplace	4.16.0	True	False	False	37m
monitoring	4.16.0	True	False	False	29m
network	4.16.0	True	False	False	38m
node-tuning	4.16.0	True	False	False	37m
openshift-apiserver	4.16.0	True	False	False	32m
openshift-controller-manager	4.16.0	True	False	False	30m
openshift-samples	4.16.0	True	False	False	32m
operator-lifecycle-manager	4.16.0	True	False	False	37m
operator-lifecycle-manager-catalog	4.16.0	True	False	False	37m
operator-lifecycle-manager-packageserver	4.16.0	True	False	False	32m
service-ca	4.16.0	True	False	False	38m
storage	4.16.0	True	False	False	37m

2. Configure the Operators that are not available.

18.3.16.1. Disabling the default OperatorHub catalog sources

Operator catalogs that source content provided by Red Hat and community projects are configured for OperatorHub by default during an OpenShift Container Platform installation. In a restricted network environment, you must disable the default catalogs as a cluster administrator.

Procedure

- Disable the sources for the default catalogs by adding **disableAllDefaultSources: true** to the **OperatorHub** object:

```
$ oc patch OperatorHub cluster --type json \
  -p '[{"op": "add", "path": "/spec/disableAllDefaultSources", "value": true}]'
```

TIP

Alternatively, you can use the web console to manage catalog sources. From the **Administration** → **Cluster Settings** → **Configuration** → **OperatorHub** page, click the **Sources** tab, where you can create, update, delete, disable, and enable individual sources.

18.3.16.2. Image registry storage configuration

The Image Registry Operator is not initially available for platforms that do not provide default storage. After installation, you must configure your registry to use storage so that the Registry Operator is made available.

Instructions are shown for configuring a persistent volume, which is required for production clusters. Where applicable, instructions are shown for configuring an empty directory as the storage location, which is available for only non-production clusters.

Additional instructions are provided for allowing the image registry to use block storage types by using the **Recreate** rollout strategy during upgrades.

18.3.16.2.1. Configuring registry storage for IBM Z

As a cluster administrator, following installation you must configure your registry to use storage.

Prerequisites

- You have access to the cluster as a user with the **cluster-admin** role.
- You have a cluster on IBM Z®.
- You have provisioned persistent storage for your cluster, such as Red Hat OpenShift Data Foundation.




IMPORTANT

OpenShift Container Platform supports **ReadWriteOnce** access for image registry storage when you have only one replica. **ReadWriteOnce** access also requires that the registry uses the **Recreate** rollout strategy. To deploy an image registry that supports high availability with two or more replicas, **ReadWriteMany** access is required.

- Must have 100Gi capacity.

Procedure

1. To configure your registry to use storage, change the **spec.storage.pvc** in the **configs.imageregistry/cluster** resource.



NOTE

When you use shared storage, review your security settings to prevent outside access.

2. Verify that you do not have a registry pod:

```
$ oc get pod -n openshift-image-registry -l docker-registry=default
```

Example output

```
No resources found in openshift-image-registry namespace
```



NOTE

If you do have a registry pod in your output, you do not need to continue with this procedure.

3. Check the registry configuration:

```
$ oc edit configs.imageregistry.operator.openshift.io
```

Example output

```
storage:
  pvc:
    claim:
```

Leave the **claim** field blank to allow the automatic creation of an **image-registry-storage** PVC.

4. Check the **clusteroperator** status:

```
$ oc get clusteroperator image-registry
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
image-registry	4.16	True	False	False	6h50m

5. Ensure that your registry is set to managed to enable building and pushing of images.

- Run:

```
$ oc edit configs.imageregistry/cluster
```

Then, change the line

```
managementState: Removed
```

to

```
managementState: Managed
```

18.3.16.2.2. Configuring storage for the image registry in non-production clusters

You must configure storage for the Image Registry Operator. For non-production clusters, you can set the image registry to an empty directory. If you do so, all images are lost if you restart the registry.

Procedure

- To set the image registry storage to an empty directory:

```
$ oc patch configs.imageregistry.operator.openshift.io cluster --type merge --patch '{"spec": {"storage":{"emptyDir":{}}}'
```



WARNING

Configure this option for only non-production clusters.

If you run this command before the Image Registry Operator initializes its components, the **oc patch** command fails with the following error:

```
Error from server (NotFound): configs.imageregistry.operator.openshift.io "cluster" not found
```

Wait a few minutes and run the command again.

18.3.17. Completing installation on user-provisioned infrastructure

After you complete the Operator configuration, you can finish installing the cluster on infrastructure that you provide.

Prerequisites

- Your control plane has initialized.
- You have completed the initial Operator configuration.

Procedure

1. Confirm that all the cluster components are online with the following command:

```
$ watch -n5 oc get clusteroperators
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.16.0	True	False	False	19m
baremetal	4.16.0	True	False	False	37m
cloud-credential	4.16.0	True	False	False	40m
cluster-autoscaler	4.16.0	True	False	False	37m
config-operator	4.16.0	True	False	False	38m
console	4.16.0	True	False	False	26m
csi-snapshot-controller	4.16.0	True	False	False	37m
dns	4.16.0	True	False	False	37m
etcd	4.16.0	True	False	False	36m
image-registry	4.16.0	True	False	False	31m
ingress	4.16.0	True	False	False	30m
insights	4.16.0	True	False	False	31m
kube-apiserver	4.16.0	True	False	False	26m
kube-controller-manager	4.16.0	True	False	False	36m
kube-scheduler	4.16.0	True	False	False	36m
kube-storage-version-migrator	4.16.0	True	False	False	37m
machine-api	4.16.0	True	False	False	29m
machine-approver	4.16.0	True	False	False	37m
machine-config	4.16.0	True	False	False	36m
marketplace	4.16.0	True	False	False	37m
monitoring	4.16.0	True	False	False	29m
network	4.16.0	True	False	False	38m
node-tuning	4.16.0	True	False	False	37m
openshift-apiserver	4.16.0	True	False	False	32m
openshift-controller-manager	4.16.0	True	False	False	30m
openshift-samples	4.16.0	True	False	False	32m
operator-lifecycle-manager	4.16.0	True	False	False	37m
operator-lifecycle-manager-catalog	4.16.0	True	False	False	37m
operator-lifecycle-manager-packageserver	4.16.0	True	False	False	32m
service-ca	4.16.0	True	False	False	38m
storage	4.16.0	True	False	False	37m

Alternatively, the following command notifies you when all of the clusters are available. It also retrieves and displays credentials:

```
$ ./openshift-install --dir <installation_directory> wait-for install-complete 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

Example output

```
INFO Waiting up to 30m0s for the cluster to initialize...
```

The command succeeds when the Cluster Version Operator finishes deploying the OpenShift Container Platform cluster from Kubernetes API server.



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

2. Confirm that the Kubernetes API server is communicating with the pods.

- a. To view a list of all pods, use the following command:

```
$ oc get pods --all-namespaces
```

Example output

```

NAMESPACE           NAME                                     READY  STATUS
RESTARTS  AGE
openshift-apiserver-operator  openshift-apiserver-operator-85cb746d55-zqhs8  1/1
Running   1      9m
openshift-apiserver          apiserver-67b9g                                1/1  Running  0
3m
openshift-apiserver          apiserver-ljcmx                                1/1  Running  0
1m
openshift-apiserver          apiserver-z25h4                                1/1  Running  0
2m
openshift-authentication-operator  authentication-operator-69d5d8bf84-vh2n8  1/1
Running   0      5m
...

```

- b. View the logs for a pod that is listed in the output of the previous command by using the following command:

```
$ oc logs <pod_name> -n <namespace> ❶
```

- ❶ Specify the pod name and namespace, as shown in the output of the previous command.

If the pod logs display, the Kubernetes API server can communicate with the cluster machines.

3. For an installation with Fibre Channel Protocol (FCP), additional steps are required to enable multipathing. Do not enable multipathing during installation. See "Enabling multipathing with kernel arguments on RHCOS" in the *Postinstallation machine configuration tasks* documentation for more information.

4. Register your cluster on the [Cluster registration](#) page.

Verification

If you have enabled secure boot during the OpenShift Container Platform bootstrap process, the following verification steps are required:

1. Debug the node by running the following command:

```
$ oc debug node/<node_name>
chroot /host
```

2. Confirm that secure boot is enabled by running the following command:

```
$ cat /sys/firmware/ipl/secure
```

Example output

```
1 1
```

- 1 The value is **1** if secure boot is enabled and **0** if secure boot is not enabled.

Additional resources

- [How to generate SOSREPORT within OpenShift Container Platform version 4 nodes without SSH.](#)

18.3.18. Next steps

- [Customize your cluster.](#)
- If the mirror registry that you used to install your cluster has a trusted CA, add it to the cluster by [configuring additional trust stores.](#)
- If necessary, you can [opt out of remote health reporting](#) .
- If necessary, see [Registering your disconnected cluster](#)

18.4. INSTALLING A CLUSTER WITH RHEL KVM ON IBM Z AND IBM LINUXONE

In OpenShift Container Platform version 4.16, you can install a cluster on IBM Z[®] or IBM[®] LinuxONE infrastructure that you provision.



NOTE

While this document refers only to IBM Z[®], all information in it also applies to IBM[®] LinuxONE.




IMPORTANT

Additional considerations exist for non-bare metal platforms. Review the information in the [guidelines for deploying OpenShift Container Platform on non-tested platforms](#) before you install an OpenShift Container Platform cluster.

18.4.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- Before you begin the installation process, you must clean the installation directory. This ensures that the required installation files are created and updated during the installation process.
- You provisioned [persistent storage using OpenShift Data Foundation](#) or other supported storage protocols for your cluster. To deploy a private image registry, you must set up persistent storage with **ReadWriteMany** access.
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.



NOTE

Be sure to also review this site list if you are configuring a proxy.

- You provisioned a RHEL Kernel Virtual Machine (KVM) system that is hosted on the logical partition (LPAR) and based on RHEL 8.6 or later. See [Red Hat Enterprise Linux 8 and 9 Life Cycle](#).

18.4.2. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.



IMPORTANT

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

18.4.3. Machine requirements for a cluster with user-provisioned infrastructure

For a cluster that contains user-provisioned infrastructure, you must deploy all of the required machines.

One or more KVM host machines based on RHEL 8.6 or later. Each RHEL KVM host machine must have libvirt installed and running. The virtual machines are provisioned under each RHEL KVM host machine.

18.4.3.1. Required machines

The smallest OpenShift Container Platform clusters require the following hosts:

Table 18.40. Minimum required hosts

Hosts	Description
One temporary bootstrap machine	The cluster requires the bootstrap machine to deploy the OpenShift Container Platform cluster on the three control plane machines. You can remove the bootstrap machine after you install the cluster.
Three control plane machines	The control plane machines run the Kubernetes and OpenShift Container Platform services that form the control plane.
At least two compute machines, which are also known as worker machines.	The workloads requested by OpenShift Container Platform users run on the compute machines.



IMPORTANT

To improve high availability of your cluster, distribute the control plane machines over different RHEL instances on at least two physical machines.

The bootstrap, control plane, and compute machines must use Red Hat Enterprise Linux CoreOS (RHCOS) as the operating system.

See [Red Hat Enterprise Linux technology capabilities and limits](#) .

18.4.3.2. Network connectivity requirements

The OpenShift Container Platform installer creates the Ignition files, which are necessary for all the Red Hat Enterprise Linux CoreOS (RHCOS) virtual machines. The automated installation of OpenShift Container Platform is performed by the bootstrap machine. It starts the installation of OpenShift Container Platform on each node, starts the Kubernetes cluster, and then finishes. During this bootstrap, the virtual machine must have an established network connection either through a Dynamic Host Configuration Protocol (DHCP) server or static IP address.

18.4.3.3. IBM Z network connectivity requirements

To install on IBM Z[®] under RHEL KVM, you need:

- A RHEL KVM host configured with an OSA or RoCE network adapter.

- Either a RHEL KVM host that is configured to use bridged networking in libvirt or MacVTap to connect the network to the guests.
See [Types of virtual network connections](#).

18.4.3.4. Host machine resource requirements

The RHEL KVM host in your environment must meet the following requirements to host the virtual machines that you plan for the OpenShift Container Platform environment. See [Getting started with virtualization](#).

You can install OpenShift Container Platform version 4.16 on the following IBM® hardware:

- IBM® z16 (all models), IBM® z15 (all models), IBM® z14 (all models)
- IBM® LinuxONE 4 (all models), IBM® LinuxONE III (all models), IBM® LinuxONE Emperor II, IBM® LinuxONE Rockhopper II

18.4.3.5. Minimum IBM Z system environment

Hardware requirements

- The equivalent of six Integrated Facilities for Linux (IFL), which are SMT2 enabled, for each cluster.
- At least one network connection to both connect to the **LoadBalancer** service and to serve data for traffic outside the cluster.



NOTE

You can use dedicated or shared IFLs to assign sufficient compute resources. Resource sharing is one of the key strengths of IBM Z®. However, you must adjust capacity correctly on each hypervisor layer and ensure sufficient resources for every OpenShift Container Platform cluster.



IMPORTANT

Since the overall performance of the cluster can be impacted, the LPARs that are used to set up the OpenShift Container Platform clusters must provide sufficient compute capacity. In this context, LPAR weight management, entitlements, and CPU shares on the hypervisor level play an important role.

Operating system requirements

- One LPAR running on RHEL 8.6 or later with KVM, which is managed by libvirt

On your RHEL KVM host, set up:

- Three guest virtual machines for OpenShift Container Platform control plane machines
- Two guest virtual machines for OpenShift Container Platform compute machines
- One guest virtual machine for the temporary OpenShift Container Platform bootstrap machine

18.4.3.6. Minimum resource requirements

Each cluster virtual machine must meet the following minimum requirements:

Virtual Machine	Operating System	vCPU [1]	Virtual RAM	Storage	IOPS
Bootstrap	RHCOS	4	16 GB	100 GB	N/A
Control plane	RHCOS	4	16 GB	100 GB	N/A
Compute	RHCOS	2	8 GB	100 GB	N/A

1. One physical core (IFL) provides two logical cores (threads) when SMT-2 is enabled. The hypervisor can provide two or more vCPUs.

18.4.3.7. Preferred IBM Z system environment

Hardware requirements

- Three LPARS that each have the equivalent of six IFLs, which are SMT2 enabled, for each cluster.
- Two network connections to both connect to the **LoadBalancer** service and to serve data for traffic outside the cluster.

Operating system requirements

- For high availability, two or three LPARs running on RHEL 8.6 or later with KVM, which are managed by libvirt.

On your RHEL KVM host, set up:

- Three guest virtual machines for OpenShift Container Platform control plane machines, distributed across the RHEL KVM host machines.
- At least six guest virtual machines for OpenShift Container Platform compute machines, distributed across the RHEL KVM host machines.
- One guest virtual machine for the temporary OpenShift Container Platform bootstrap machine.
- To ensure the availability of integral components in an overcommitted environment, increase the priority of the control plane by using **cpu_shares**. Do the same for infrastructure nodes, if they exist. See [schedinfo](#) in IBM® Documentation.

18.4.3.8. Preferred resource requirements

The preferred requirements for each cluster virtual machine are:

Virtual Machine	Operating System	vCPU	Virtual RAM	Storage
Bootstrap	RHCOS	4	16 GB	120 GB
Control plane	RHCOS	8	16 GB	120 GB

Virtual Machine	Operating System	vCPU	Virtual RAM	Storage
Compute	RHCOS	6	8 GB	120 GB

18.4.3.9. Certificate signing requests management

Because your cluster has limited access to automatic machine management when you use infrastructure that you provision, you must provide a mechanism for approving cluster certificate signing requests (CSRs) after installation. The **kube-controller-manager** only approves the kubelet client CSRs. The **machine-approver** cannot guarantee the validity of a serving certificate that is requested by using kubelet credentials because it cannot confirm that the correct machine issued the request. You must determine and implement a method of verifying the validity of the kubelet serving certificate requests and approving them.

Additional resources

- [Recommended host practices for IBM Z® & IBM® LinuxONE environments](#)

18.4.3.10. Networking requirements for user-provisioned infrastructure

All the Red Hat Enterprise Linux CoreOS (RHCOS) machines require networking to be configured in **initramfs** during boot to fetch their Ignition config files.

During the initial boot, the machines require an IP address configuration that is set either through a DHCP server or statically by providing the required boot options. After a network connection is established, the machines download their Ignition config files from an HTTP or HTTPS server. The Ignition config files are then used to set the exact state of each machine. The Machine Config Operator completes more changes to the machines, such as the application of new certificates or keys, after installation.

It is recommended to use a DHCP server for long-term management of the cluster machines. Ensure that the DHCP server is configured to provide persistent IP addresses, DNS server information, and hostnames to the cluster machines.



NOTE

If a DHCP service is not available for your user-provisioned infrastructure, you can instead provide the IP networking configuration and the address of the DNS server to the nodes at RHCOS install time. These can be passed as boot arguments if you are installing from an ISO image. See the *Installing RHCOS and starting the OpenShift Container Platform bootstrap process* section for more information about static IP provisioning and advanced networking options.

The Kubernetes API server must be able to resolve the node names of the cluster machines. If the API servers and worker nodes are in different zones, you can configure a default DNS search zone to allow the API server to resolve the node names. Another supported approach is to always refer to hosts by their fully-qualified domain names in both the node objects and all DNS requests.

18.4.3.10.1. Setting the cluster node hostnames through DHCP

On Red Hat Enterprise Linux CoreOS (RHCOS) machines, the hostname is set through NetworkManager. By default, the machines obtain their hostname through DHCP. If the hostname is not provided by DHCP, set statically through kernel arguments, or another method, it is obtained through a


reverse DNS lookup. Reverse DNS lookup occurs after the network has been initialized on a node and can take time to resolve. Other system services can start prior to this and detect the hostname as **localhost** or similar. You can avoid this by using DHCP to provide the hostname for each cluster node.

Additionally, setting the hostnames through DHCP can bypass any manual DNS record name configuration errors in environments that have a DNS split-horizon implementation.

18.4.3.10.2. Network connectivity requirements


You must configure the network connectivity between machines to allow OpenShift Container Platform cluster components to communicate. Each machine must be able to resolve the hostnames of all other machines in the cluster.

This section provides details about the ports that are required.



IMPORTANT

In connected OpenShift Container Platform environments, all nodes are required to have internet access to pull images for platform containers and provide telemetry data to Red Hat.



NOTE

The RHEL KVM host must be configured to use bridged networking in libvirt or MacVTap to connect the network to the virtual machines. The virtual machines must have access to the network, which is attached to the RHEL KVM host. Virtual Networks, for example network address translation (NAT), within KVM are not a supported configuration.

Table 18.41. Ports used for all-machine to all-machine communications

Protocol	Port	Description
ICMP	N/A	Network reachability tests
TCP	1936	Metrics
	9000-9999	Host level services, including the node exporter on ports 9100-9101 and the Cluster Version Operator on port 9099 .
	10250-10259	The default ports that Kubernetes reserves
UDP	4789	VXLAN
	6081	Geneve
	9000-9999	Host level services, including the node exporter on ports 9100-9101 .
	500	IPsec IKE packets
	4500	IPsec NAT-T packets

Protocol	Port	Description
	123	Network Time Protocol (NTP) on UDP port 123 If an external NTP time server is configured, you must open UDP port 123 .
TCP/UDP	30000-32767	Kubernetes node port
ESP	N/A	IPsec Encapsulating Security Payload (ESP)

Table 18.42. Ports used for all-machine to control plane communications

Protocol	Port	Description
TCP	6443	Kubernetes API

Table 18.43. Ports used for control plane machine to control plane machine communications

Protocol	Port	Description
TCP	2379-2380	etcd server and peer ports

NTP configuration for user-provisioned infrastructure

OpenShift Container Platform clusters are configured to use a public Network Time Protocol (NTP) server by default. If you want to use a local enterprise NTP server, or if your cluster is being deployed in a disconnected network, you can configure the cluster to use a specific time server. For more information, see the documentation for *Configuring chrony time service*.

If a DHCP server provides NTP server information, the chrony time service on the Red Hat Enterprise Linux CoreOS (RHCOS) machines read the information and can sync the clock with the NTP servers.

Additional resources

- [Configuring chrony time service](#)

18.4.3.11. User-provisioned DNS requirements

In OpenShift Container Platform deployments, DNS name resolution is required for the following components:


- The Kubernetes API
- The OpenShift Container Platform application wildcard
- The bootstrap, control plane, and compute machines

Reverse DNS resolution is also required for the Kubernetes API, the bootstrap machine, the control plane machines, and the compute machines.

DNS A/AAAA or CNAME records are used for name resolution and PTR records are used for reverse name resolution. The reverse records are important because Red Hat Enterprise Linux CoreOS (RHCOS) uses the reverse records to set the hostnames for all the nodes, unless the hostnames are provided by DHCP. Additionally, the reverse records are used to generate the certificate signing requests (CSR) that OpenShift Container Platform needs to operate.

The following DNS records are required for a user-provisioned OpenShift Container Platform cluster and they must be in place before installation. In each record, **<cluster_name>** is the cluster name and **<base_domain>** is the base domain that you specify in the **install-config.yaml** file. A complete DNS record takes the form: **<component>.<cluster_name>.<base_domain>**.

Table 18.44. Required DNS records

Component	Record	Description
Kubernetes API	api.<cluster_name>.<base_domain>	A DNS A/AAAA or CNAME record, and a DNS PTR record, to identify the API load balancer. These records must be resolvable by both clients external to the cluster and from all the nodes within the cluster.
	api-int.<cluster_name>.<base_domain>	A DNS A/AAAA or CNAME record, and a DNS PTR record, to internally identify the API load balancer. These records must be resolvable from all the nodes within the cluster.
		 <p>IMPORTANT</p> <p>The API server must be able to resolve the worker nodes by the hostnames that are recorded in Kubernetes. If the API server cannot resolve the node names, then proxied API calls can fail, and you cannot retrieve logs from pods.</p>
Routes	*.apps.<cluster_name>.<base_domain>	A wildcard DNS A/AAAA or CNAME record that refers to the application ingress load balancer. The application ingress load balancer targets the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default. These records must be resolvable by both clients external to the cluster and from all the nodes within the cluster.
		For example, console-openshift-console.apps.<cluster_name>.<base_domain> is used as a wildcard route to the OpenShift Container Platform console.
Bootstrap machine	bootstrap.<cluster_name>.<base_domain>	A DNS A/AAAA or CNAME record, and a DNS PTR record, to identify the bootstrap machine. These records must be resolvable by the nodes within the cluster.
Control plane machines	<control_plane><n>.<cluster_name>.<base_domain>	DNS A/AAAA or CNAME records and DNS PTR records to identify each machine for the control plane nodes. These records must be resolvable by the nodes within the cluster.

Component	Record	Description
Compute machines	<compute><n>. <cluster_name>. <base_domain>.	DNS A/AAAA or CNAME records and DNS PTR records to identify each machine for the worker nodes. These records must be resolvable by the nodes within the cluster.

**NOTE**

In OpenShift Container Platform 4.4 and later, you do not need to specify etcd host and SRV records in your DNS configuration.

TIP

You can use the **dig** command to verify name and reverse name resolution. See the section on *Validating DNS resolution for user-provisioned infrastructure* for detailed validation steps.

18.4.3.11.1. Example DNS configuration for user-provisioned clusters

This section provides A and PTR record configuration samples that meet the DNS requirements for deploying OpenShift Container Platform on user-provisioned infrastructure. The samples are not meant to provide advice for choosing one DNS solution over another.

In the examples, the cluster name is **ocp4** and the base domain is **example.com**.

Example DNS A record configuration for a user-provisioned cluster

The following example is a BIND zone file that shows sample A records for name resolution in a user-provisioned cluster.

Example 18.7. Sample DNS zone database

```
$TTL 1W
@ IN SOA ns1.example.com. root (
    2019070700 ; serial
    3H ; refresh (3 hours)
    30M ; retry (30 minutes)
    2W ; expiry (2 weeks)
    1W ) ; minimum (1 week)
IN NS ns1.example.com.
IN MX 10 smtp.example.com.
;
;
ns1.example.com. IN A 192.168.1.5
smtp.example.com. IN A 192.168.1.5
;
helper.example.com. IN A 192.168.1.5
```

```

helper.ocp4.example.com. IN A 192.168.1.5
;
api.ocp4.example.com. IN A 192.168.1.5 1
api-int.ocp4.example.com. IN A 192.168.1.5 2
;
*.apps.ocp4.example.com. IN A 192.168.1.5 3
;
bootstrap.ocp4.example.com. IN A 192.168.1.96 4
;
control-plane0.ocp4.example.com. IN A 192.168.1.97 5
control-plane1.ocp4.example.com. IN A 192.168.1.98 6
control-plane2.ocp4.example.com. IN A 192.168.1.99 7
;
compute0.ocp4.example.com. IN A 192.168.1.11 8
compute1.ocp4.example.com. IN A 192.168.1.7 9
;
;EOF

```

- 1 Provides name resolution for the Kubernetes API. The record refers to the IP address of the API load balancer.
- 2 Provides name resolution for the Kubernetes API. The record refers to the IP address of the API load balancer and is used for internal cluster communications.
- 3 Provides name resolution for the wildcard routes. The record refers to the IP address of the application ingress load balancer. The application ingress load balancer targets the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.



NOTE

In the example, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

- 4 Provides name resolution for the bootstrap machine.
- 5 6 7 Provides name resolution for the control plane machines.
- 8 9 Provides name resolution for the compute machines.

Example DNS PTR record configuration for a user-provisioned cluster

The following example BIND zone file shows sample PTR records for reverse name resolution in a user-provisioned cluster.

Example 18.8. Sample DNS zone database for reverse records

```

$TTL 1W
@ IN SOA ns1.example.com. root (
  2019070700 ; serial

```

```

3H ; refresh (3 hours)
30M ; retry (30 minutes)
2W ; expiry (2 weeks)
1W ) ; minimum (1 week)
IN NS ns1.example.com.
;
5.1.168.192.in-addr.arpa. IN PTR api.ocp4.example.com. 1
5.1.168.192.in-addr.arpa. IN PTR api-int.ocp4.example.com. 2
;
96.1.168.192.in-addr.arpa. IN PTR bootstrap.ocp4.example.com. 3
;
97.1.168.192.in-addr.arpa. IN PTR control-plane0.ocp4.example.com. 4
98.1.168.192.in-addr.arpa. IN PTR control-plane1.ocp4.example.com. 5
99.1.168.192.in-addr.arpa. IN PTR control-plane2.ocp4.example.com. 6
;
11.1.168.192.in-addr.arpa. IN PTR compute0.ocp4.example.com. 7
7.1.168.192.in-addr.arpa. IN PTR compute1.ocp4.example.com. 8
;
;EOF

```

- 1 Provides reverse DNS resolution for the Kubernetes API. The PTR record refers to the record name of the API load balancer.
- 2 Provides reverse DNS resolution for the Kubernetes API. The PTR record refers to the record name of the API load balancer and is used for internal cluster communications.
- 3 Provides reverse DNS resolution for the bootstrap machine.
- 4 5 6 Provides reverse DNS resolution for the control plane machines.
- 7 8 Provides reverse DNS resolution for the compute machines.



NOTE

A PTR record is not required for the OpenShift Container Platform application wildcard.

18.4.3.12. Load balancing requirements for user-provisioned infrastructure

Before you install OpenShift Container Platform, you must provision the API and application Ingress load balancing infrastructure. In production scenarios, you can deploy the API and application Ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.




NOTE

If you want to deploy the API and application Ingress load balancers with a Red Hat Enterprise Linux (RHEL) instance, you must purchase the RHEL subscription separately.

The load balancing infrastructure must meet the following requirements:

1. **API load balancer.** Provides a common endpoint for users, both human and machine, to interact with and configure the platform. Configure the following conditions:

- Layer 4 load balancing only. This can be referred to as Raw TCP or SSL Passthrough mode.
- A stateless load balancing algorithm. The options vary based on the load balancer implementation.



IMPORTANT

Do not configure session persistence for an API load balancer. Configuring session persistence for a Kubernetes API server might cause performance issues from excess application traffic for your OpenShift Container Platform cluster and the Kubernetes API that runs inside the cluster.

Configure the following ports on both the front and back of the load balancers:

Table 18.45. API load balancer

Port	Back-end machines (pool members)	Internal	External	Description
6443	Bootstrap and control plane. You remove the bootstrap machine from the load balancer after the bootstrap machine initializes the cluster control plane. You must configure the /readyz endpoint for the API server health check probe.	X	X	Kubernetes API server
22623	Bootstrap and control plane. You remove the bootstrap machine from the load balancer after the bootstrap machine initializes the cluster control plane.	X		Machine config server



NOTE

The load balancer must be configured to take a maximum of 30 seconds from the time the API server turns off the **/readyz** endpoint to the removal of the API server instance from the pool. Within the time frame after **/readyz** returns an error or becomes healthy, the endpoint must have been removed or added. Probing every 5 or 10 seconds, with two successful requests to become healthy and three to become unhealthy, are well-tested values.

2. **Application Ingress load balancer.** Provides an ingress point for application traffic flowing in from outside the cluster. A working configuration for the Ingress router is required for an OpenShift Container Platform cluster.

Configure the following conditions:

- Layer 4 load balancing only. This can be referred to as Raw TCP or SSL Passthrough mode.
- A connection-based or session-based persistence is recommended, based on the options available and types of applications that will be hosted on the platform.

TIP

If the true IP address of the client can be seen by the application Ingress load balancer, enabling source IP-based session persistence can improve performance for applications that use end-to-end TLS encryption.

Configure the following ports on both the front and back of the load balancers:

Table 18.46. Application Ingress load balancer

Port	Back-end machines (pool members)	Internal	External	Description
443	The machines that run the Ingress Controller pods, compute, or worker, by default.	X	X	HTTPS traffic
80	The machines that run the Ingress Controller pods, compute, or worker, by default.	X	X	HTTP traffic

**NOTE**

If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application Ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes.

18.4.3.12.1. Example load balancer configuration for user-provisioned clusters

This section provides an example API and application Ingress load balancer configuration that meets the load balancing requirements for user-provisioned clusters. The sample is an `/etc/haproxy/haproxy.cfg` configuration for an HAProxy load balancer. The example is not meant to provide advice for choosing one load balancing solution over another.

In the example, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

**NOTE**

If you are using HAProxy as a load balancer and SELinux is set to **enforcing**, you must ensure that the HAProxy service can bind to the configured TCP port by running **setsebool -P haproxy_connect_any=1**.

Example 18.9. Sample API and application Ingress load balancer configuration

```
global
log      127.0.0.1 local2
pidfile  /var/run/haproxy.pid
maxconn 4000
daemon
defaults
```

```

mode                http
log                 global
option             dontlognull
option http-server-close
option             redispatch
retries            3
timeout http-request 10s
timeout queue       1m
timeout connect     10s
timeout client      1m
timeout server      1m
timeout http-keep-alive 10s
timeout check       10s
maxconn            3000
listen api-server-6443 1
bind *:6443
mode tcp
option httpchk GET /readyz HTTP/1.0
option log-health-checks
balance roundrobin
server bootstrap bootstrap.ocp4.example.com:6443 verify none check check-ssl inter 10s fall 2
rise 3 backup 2
server master0 master0.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
fall 2 rise 3
server master1 master1.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
fall 2 rise 3
server master2 master2.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
fall 2 rise 3
listen machine-config-server-22623 3
bind *:22623
mode tcp
server bootstrap bootstrap.ocp4.example.com:22623 check inter 1s backup 4
server master0 master0.ocp4.example.com:22623 check inter 1s
server master1 master1.ocp4.example.com:22623 check inter 1s
server master2 master2.ocp4.example.com:22623 check inter 1s
listen ingress-router-443 5
bind *:443
mode tcp
balance source
server compute0 compute0.ocp4.example.com:443 check inter 1s
server compute1 compute1.ocp4.example.com:443 check inter 1s
listen ingress-router-80 6
bind *:80
mode tcp
balance source
server compute0 compute0.ocp4.example.com:80 check inter 1s
server compute1 compute1.ocp4.example.com:80 check inter 1s

```

- 1** Port **6443** handles the Kubernetes API traffic and points to the control plane machines.
- 2** **4** The bootstrap entries must be in place before the OpenShift Container Platform cluster installation and they must be removed after the bootstrap process is complete.
- 3** Port **22623** handles the machine config server traffic and points to the control plane machines.

- 5 Port **443** handles the HTTPS traffic and points to the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.
- 6 Port **80** handles the HTTP traffic and points to the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.



NOTE

If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application Ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes.

TIP

If you are using HAProxy as a load balancer, you can check that the **haproxy** process is listening on ports **6443**, **22623**, **443**, and **80** by running **netstat -nltp** on the HAProxy node.

18.4.4. Preparing the user-provisioned infrastructure

Before you install OpenShift Container Platform on user-provisioned infrastructure, you must prepare the underlying infrastructure.

This section provides details about the high-level steps required to set up your cluster infrastructure in preparation for an OpenShift Container Platform installation. This includes configuring IP networking and network connectivity for your cluster nodes, enabling the required ports through your firewall, and setting up the required DNS and load balancing infrastructure.

After preparation, your cluster infrastructure must meet the requirements outlined in the *Requirements for a cluster with user-provisioned infrastructure* section.

Prerequisites

- You have reviewed the [OpenShift Container Platform 4.x Tested Integrations](#) page.
- You have reviewed the infrastructure requirements detailed in the *Requirements for a cluster with user-provisioned infrastructure* section.

Procedure

1. If you are using DHCP to provide the IP networking configuration to your cluster nodes, configure your DHCP service.
 - a. Add persistent IP addresses for the nodes to your DHCP server configuration. In your configuration, match the MAC address of the relevant network interface to the intended IP address for each node.
 - b. When you use DHCP to configure IP addressing for the cluster machines, the machines also obtain the DNS server information through DHCP. Define the persistent DNS server address that is used by the cluster nodes through your DHCP server configuration.

**NOTE**

If you are not using a DHCP service, you must provide the IP networking configuration and the address of the DNS server to the nodes at RHCOS install time. These can be passed as boot arguments if you are installing from an ISO image. See the *Installing RHCOS and starting the OpenShift Container Platform bootstrap process* section for more information about static IP provisioning and advanced networking options.

- c. Define the hostnames of your cluster nodes in your DHCP server configuration. See the *Setting the cluster node hostnames through DHCP* section for details about hostname considerations.

**NOTE**

If you are not using a DHCP service, the cluster nodes obtain their hostname through a reverse DNS lookup.

2. Choose to perform either a fast track installation of Red Hat Enterprise Linux CoreOS (RHCOS) or a full installation of Red Hat Enterprise Linux CoreOS (RHCOS). For the full installation, you must set up an HTTP or HTTPS server to provide Ignition files and install images to the cluster nodes. For the fast track installation an HTTP or HTTPS server is not required, however, a DHCP server is required. See sections "Fast-track installation: Creating Red Hat Enterprise Linux CoreOS (RHCOS) machines" and "Full installation: Creating Red Hat Enterprise Linux CoreOS (RHCOS) machines".
3. Ensure that your network infrastructure provides the required network connectivity between the cluster components. See the *Networking requirements for user-provisioned infrastructure* section for details about the requirements.
4. Configure your firewall to enable the ports required for the OpenShift Container Platform cluster components to communicate. See *Networking requirements for user-provisioned infrastructure* section for details about the ports that are required.

**IMPORTANT**

By default, port **1936** is accessible for an OpenShift Container Platform cluster, because each control plane node needs access to this port.

Avoid using the Ingress load balancer to expose this port, because doing so might result in the exposure of sensitive information, such as statistics and metrics, related to Ingress Controllers.

5. Setup the required DNS infrastructure for your cluster.
 - a. Configure DNS name resolution for the Kubernetes API, the application wildcard, the bootstrap machine, the control plane machines, and the compute machines.
 - b. Configure reverse DNS resolution for the Kubernetes API, the bootstrap machine, the control plane machines, and the compute machines.
See the *User-provisioned DNS requirements* section for more information about the OpenShift Container Platform DNS requirements.
6. Validate your DNS configuration.

- a. From your installation node, run DNS lookups against the record names of the Kubernetes API, the wildcard routes, and the cluster nodes. Validate that the IP addresses in the responses correspond to the correct components.
 - b. From your installation node, run reverse DNS lookups against the IP addresses of the load balancer and the cluster nodes. Validate that the record names in the responses correspond to the correct components.
See the *Validating DNS resolution for user-provisioned infrastructure* section for detailed DNS validation steps.
7. Provision the required API and application ingress load balancing infrastructure. See the *Load balancing requirements for user-provisioned infrastructure* section for more information about the requirements.




NOTE

Some load balancing solutions require the DNS name resolution for the cluster nodes to be in place before the load balancing is initialized.

18.4.5. Validating DNS resolution for user-provisioned infrastructure

You can validate your DNS configuration before installing OpenShift Container Platform on user-provisioned infrastructure.



IMPORTANT

The validation steps detailed in this section must succeed before you install your cluster.

Prerequisites

- You have configured the required DNS records for your user-provisioned infrastructure.

Procedure

1. From your installation node, run DNS lookups against the record names of the Kubernetes API, the wildcard routes, and the cluster nodes. Validate that the IP addresses contained in the responses correspond to the correct components.
 - a. Perform a lookup against the Kubernetes API record name. Check that the result points to the IP address of the API load balancer:

```
$ dig +noall +answer @<nameserver_ip> api.<cluster_name>.<base_domain> 1
```

- 1 Replace **<nameserver_ip>** with the IP address of the nameserver, **<cluster_name>** with your cluster name, and **<base_domain>** with your base domain name.

Example output

```
api.ocp4.example.com. 604800 IN A 192.168.1.5
```

- b. Perform a lookup against the Kubernetes internal API record name. Check that the result points to the IP address of the API load balancer:

```
$ dig +noall +answer @<nameserver_ip> api-int.<cluster_name>.<base_domain>
```

Example output

```
api-int.ocp4.example.com. 604800 IN A 192.168.1.5
```

- c. Test an example ***.apps.<cluster_name>.<base_domain>** DNS wildcard lookup. All of the application wildcard lookups must resolve to the IP address of the application ingress load balancer:

```
$ dig +noall +answer @<nameserver_ip> random.apps.<cluster_name>.<base_domain>
```

Example output

```
random.apps.ocp4.example.com. 604800 IN A 192.168.1.5
```



NOTE

In the example outputs, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

You can replace **random** with another wildcard value. For example, you can query the route to the OpenShift Container Platform console:

```
$ dig +noall +answer @<nameserver_ip> console-openshift-console.apps.<cluster_name>.<base_domain>
```

Example output

```
console-openshift-console.apps.ocp4.example.com. 604800 IN A 192.168.1.5
```

- d. Run a lookup against the bootstrap DNS record name. Check that the result points to the IP address of the bootstrap node:

```
$ dig +noall +answer @<nameserver_ip> bootstrap.<cluster_name>.<base_domain>
```

Example output

```
bootstrap.ocp4.example.com. 604800 IN A 192.168.1.96
```

- e. Use this method to perform lookups against the DNS record names for the control plane and compute nodes. Check that the results correspond to the IP addresses of each node.
2. From your installation node, run reverse DNS lookups against the IP addresses of the load balancer and the cluster nodes. Validate that the record names contained in the responses correspond to the correct components.
 - a. Perform a reverse lookup against the IP address of the API load balancer. Check that the response includes the record names for the Kubernetes API and the Kubernetes internal

API:

```
$ dig +noall +answer @<nameserver_ip> -x 192.168.1.5
```

Example output

```
5.1.168.192.in-addr.arpa. 604800 IN PTR api-int.ocp4.example.com. 1
5.1.168.192.in-addr.arpa. 604800 IN PTR api.ocp4.example.com. 2
```

- 1 Provides the record name for the Kubernetes internal API.
- 2 Provides the record name for the Kubernetes API.



NOTE

A PTR record is not required for the OpenShift Container Platform application wildcard. No validation step is needed for reverse DNS resolution against the IP address of the application ingress load balancer.

- b. Perform a reverse lookup against the IP address of the bootstrap node. Check that the result points to the DNS record name of the bootstrap node:

```
$ dig +noall +answer @<nameserver_ip> -x 192.168.1.96
```

Example output

```
96.1.168.192.in-addr.arpa. 604800 IN PTR bootstrap.ocp4.example.com.
```

- c. Use this method to perform reverse lookups against the IP addresses for the control plane and compute nodes. Check that the results correspond to the DNS record names of each node.

18.4.6. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.



NOTE

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.



NOTE

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```



NOTE

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1 Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

18.4.7. Obtaining the installation program

Before you install OpenShift Container Platform, download the installation file on your provisioning machine.

Prerequisites

- You have a machine that runs Linux, for example Red Hat Enterprise Linux 8, with 500 MB of local disk space.

Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select your infrastructure provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.



IMPORTANT

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.



IMPORTANT

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.


4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#) . This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

18.4.8. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.



IMPORTANT

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**.
To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.

3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.



NOTE

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your **PATH**.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

18.4.9. Manually creating the installation configuration file

Installing the cluster requires that you manually create the installation configuration file.


Prerequisites

- You have an SSH public key on your local machine to provide to the installation program. The key will be used for SSH authentication onto your cluster nodes for debugging and disaster recovery.
- You have obtained the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create an installation directory to store your required installation assets in:

```
$ mkdir <installation_directory>
```



IMPORTANT

You must create a directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.


2. Customize the sample **install-config.yaml** file template that is provided and save it in the **<installation_directory>**.



NOTE

You must name this configuration file **install-config.yaml**.

3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the next step of the installation process. You must back it up now.

Additional resources

- [Installation configuration parameters for IBM Z®](#)

18.4.9.1. Sample install-config.yaml file for IBM Z

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.

```
apiVersion: v1
baseDomain: example.com 1
compute: 2
- hyperthreading: Enabled 3
name: worker
replicas: 0 4
architecture: s390x
```



```

controlPlane: 5
  hyperthreading: Enabled 6
  name: master
  replicas: 3 7
  architecture: s390x
metadata:
  name: test 8
networking:
  clusterNetwork:
  - cidr: 10.128.0.0/14 9
    hostPrefix: 23 10
  networkType: OVNKubernetes 11
  serviceNetwork: 12
  - 172.30.0.0/16
platform:
  none: {} 13
fips: false 14
pullSecret: '{"auths": ...}' 15
sshKey: 'ssh-ed25519 AAAA...' 16

```

- 1 The base domain of the cluster. All DNS records must be sub-domains of this base and include the cluster name.
- 2 5 The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.
- 3 6 Specifies whether to enable or disable simultaneous multithreading (SMT), or hyperthreading. By default, SMT is enabled to increase the performance of the cores in your machines. You can disable it by setting the parameter value to **Disabled**. If you disable SMT, you must disable it in all cluster machines; this includes both control plane and compute machines.



NOTE

Simultaneous multithreading (SMT) is enabled by default. If SMT is not available on your OpenShift Container Platform nodes, the **hyperthreading** parameter has no effect.



IMPORTANT

If you disable **hyperthreading**, whether on your OpenShift Container Platform nodes or in the **install-config.yaml** file, ensure that your capacity planning accounts for the dramatically decreased machine performance.

- 4 You must set this value to **0** when you install OpenShift Container Platform on user-provisioned infrastructure. In installer-provisioned installations, the parameter controls the number of compute machines that the cluster creates and manages for you. In user-provisioned installations, you must manually deploy the compute machines before you finish installing the cluster.

**NOTE**

If you are installing a three-node cluster, do not deploy any compute machines when you install the Red Hat Enterprise Linux CoreOS (RHCOS) machines.

- 7 The number of control plane machines that you add to the cluster. Because the cluster uses these values as the number of etcd endpoints in the cluster, the value must match the number of control plane machines that you deploy.
- 8 The cluster name that you specified in your DNS records.
- 9 A block of IP addresses from which pod IP addresses are allocated. This block must not overlap with existing physical networks. These IP addresses are used for the pod network. If you need to access the pods from an external network, you must configure load balancers and routers to manage the traffic.

**NOTE**

Class E CIDR range is reserved for a future use. To use the Class E CIDR range, you must ensure your networking environment accepts the IP addresses within the Class E CIDR range.

- 10 The subnet prefix length to assign to each individual node. For example, if **hostPrefix** is set to **23**, then each node is assigned a **/23** subnet out of the given **cidr**, which allows for 510 ($2^{(32 - 23)} - 2$) pod IP addresses. If you are required to provide access to nodes from an external network, configure load balancers and routers to manage the traffic.
- 11 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 12 The IP address pool to use for service IP addresses. You can enter only one IP address pool. This block must not overlap with existing physical networks. If you need to access the services from an external network, configure load balancers and routers to manage the traffic.
- 13 You must set the platform to **none**. You cannot provide additional platform configuration variables for IBM Z[®] infrastructure.

**IMPORTANT**

Clusters that are installed with the platform type **none** are unable to use some features, such as managing compute machines with the Machine API. This limitation applies even if the compute machines that are attached to the cluster are installed on a platform that would normally support the feature. This parameter cannot be changed after installation.

- 14 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.



IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 15 The [pull secret from Red Hat OpenShift Cluster Manager](#) . This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.
- 16 The SSH public key for the **core** user in Red Hat Enterprise Linux CoreOS (RHCOS).



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

18.4.9.2. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```

apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5

```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.



NOTE

The installation program does not support the proxy **readinessEndpoints** field.



NOTE

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.



NOTE

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

18.4.9.3. Configuring a three-node cluster

Optionally, you can deploy zero compute machines in a minimal three node cluster that consists of three control plane machines only. This provides smaller, more resource efficient clusters for cluster administrators and developers to use for testing, development, and production.

In three-node OpenShift Container Platform environments, the three control plane machines are schedulable, which means that your application workloads are scheduled to run on them.

Prerequisites

- You have an existing **install-config.yaml** file.

Procedure

- Ensure that the number of compute replicas is set to **0** in your **install-config.yaml** file, as shown in the following **compute** stanza:

```
compute:
- name: worker
  platform: {}
  replicas: 0
```



NOTE

You must set the value of the **replicas** parameter for the compute machines to **0** when you install OpenShift Container Platform on user-provisioned infrastructure, regardless of the number of compute machines you are deploying. In installer-provisioned installations, the parameter controls the number of compute machines that the cluster creates and manages for you. This does not apply to user-provisioned installations, where the compute machines are deployed manually.



NOTE

The preferred resource for control plane nodes is six vCPUs and 21 GB. For three control plane nodes this is the memory + vCPU equivalent of a minimum five-node cluster. You should back the three nodes, each installed on a 120 GB disk, with three IFLs that are SMT2 enabled. The minimum tested setup is three vCPUs and 10 GB on a 120 GB disk for each control plane node.

For three-node cluster installations, follow these next steps:

- If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your

application ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes. See the *Load balancing requirements for user-provisioned infrastructure* section for more information.

- When you create the Kubernetes manifest files in the following procedure, ensure that the **mastersSchedulable** parameter in the `<installation_directory>/manifests/cluster-scheduler-02-config.yml` file is set to **true**. This enables your application workloads to run on the control plane nodes.
- Do not deploy any compute nodes when you create the Red Hat Enterprise Linux CoreOS (RHCOS) machines.

18.4.10. Cluster Network Operator configuration

The configuration for the cluster network is specified as part of the Cluster Network Operator (CNO) configuration and stored in a custom resource (CR) object that is named **cluster**. The CR specifies the fields for the **Network** API in the **operator.openshift.io** API group.

The CNO configuration inherits the following fields during cluster installation from the **Network** API in the **Network.config.openshift.io** API group:

clusterNetwork

IP address pools from which pod IP addresses are allocated.

serviceNetwork

IP address pool for services.

defaultNetwork.type

Cluster network plugin. **OVNKubernetes** is the only supported plugin during installation.

You can specify the cluster network plugin configuration for your cluster by setting the fields for the **defaultNetwork** object in the CNO object named **cluster**.

18.4.10.1. Cluster Network Operator configuration object

The fields for the Cluster Network Operator (CNO) are described in the following table:

Table 18.47. Cluster Network Operator configuration object


Field	Type	Description
metadata.name	string	The name of the CNO object. This name is always cluster .
spec.clusterNetwork	array	A list specifying the blocks of IP addresses from which pod IP addresses are allocated and the subnet prefix length assigned to each individual node in the cluster. For example: <pre>spec: clusterNetwork: - cidr: 10.128.0.0/19 hostPrefix: 23 - cidr: 10.128.32.0/19 hostPrefix: 23</pre>

Field	Type	Description
spec.serviceNetwork	array	<p>A block of IP addresses for services. The OpenShift SDN and OVN-Kubernetes network plugins support only a single IP address block for the service network. For example:</p> <pre>spec: serviceNetwork: - 172.30.0.0/14</pre> <p>You can customize this field only in the install-config.yaml file before you create the manifests. The value is read-only in the manifest file.</p>
spec.defaultNetwork	object	Configures the network plugin for the cluster network.
spec.kubeProxyConfig	object	The fields for this object specify the kube-proxy configuration. If you are using the OVN-Kubernetes cluster network plugin, the kube-proxy configuration has no effect.

defaultNetwork object configuration

The values for the **defaultNetwork** object are defined in the following table:

Table 18.48. **defaultNetwork** object

Field	Type	Description
type	string	<p>OVNKubernetes. The Red Hat OpenShift Networking network plugin is selected during installation. This value cannot be changed after cluster installation.</p> <div style="display: flex; align-items: flex-start;">  <div> <p>NOTE</p> <p>OpenShift Container Platform uses the OVN-Kubernetes network plugin by default. OpenShift SDN is no longer available as an installation choice for new clusters.</p> </div> </div>
ovnKubernetesConfig	object	This object is only valid for the OVN-Kubernetes network plugin.

Configuration for the OVN-Kubernetes network plugin

The following table describes the configuration fields for the OVN-Kubernetes network plugin:

Table 18.49. **ovnKubernetesConfig** object


Field	Type	Description
mtu	integer	<p>The maximum transmission unit (MTU) for the Geneve (Generic Network Virtualization Encapsulation) overlay network. This is detected automatically based on the MTU of the primary network interface. You do not normally need to override the detected MTU.</p> <p>If the auto-detected value is not what you expect it to be, confirm that the MTU on the primary network interface on your nodes is correct. You cannot use this option to change the MTU value of the primary network interface on the nodes.</p> <p>If your cluster requires different MTU values for different nodes, you must set this value to 100 less than the lowest MTU value in your cluster. For example, if some nodes in your cluster have an MTU of 9001, and some have an MTU of 1500, you must set this value to 1400.</p>
genevePort	integer	The port to use for all Geneve packets. The default value is 6081 . This value cannot be changed after cluster installation.
ipsecConfig	object	Specify a configuration object for customizing the IPsec configuration.
ipv4	object	Specifies a configuration object for IPv4 settings.
ipv6	object	Specifies a configuration object for IPv6 settings.
policyAuditConfig	object	Specify a configuration object for customizing network policy audit logging. If unset, the defaults audit log settings are used.
gatewayConfig	object	<p>Optional: Specify a configuration object for customizing how egress traffic is sent to the node gateway.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p>NOTE</p> <p>While migrating egress traffic, you can expect some disruption to workloads and service traffic until the Cluster Network Operator (CNO) successfully rolls out the changes.</p> </div> </div>

Table 18.50. `ovnKubernetesConfig.ipv4` object

Field	Type	Description
-------	------	-------------

Field	Type	Description
internalTransitSwitchSubnet	string	<p>If your existing network infrastructure overlaps with the 100.88.0.0/16 IPv4 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. The subnet for the distributed transit switch that enables east-west traffic. This subnet cannot overlap with any other subnets used by OVN-Kubernetes or on the host itself. It must be large enough to accommodate one IP address per node in your cluster.</p> <p>The default value is 100.88.0.0/16.</p>
internalJoinSubnet	string	<p>If your existing network infrastructure overlaps with the 100.64.0.0/16 IPv4 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. You must ensure that the IP address range does not overlap with any other subnet used by your OpenShift Container Platform installation. The IP address range must be larger than the maximum number of nodes that can be added to the cluster. For example, if the clusterNetwork.cidr value is 10.128.0.0/14 and the clusterNetwork.hostPrefix value is /23, then the maximum number of nodes is 2⁽²³⁻¹⁴⁾=512.</p> <p>The default value is 100.64.0.0/16.</p>

Table 18.51. `ovnKubernetesConfig.ipv6` object

Field	Type	Description
internalTransitSwitchSubnet	string	<p>If your existing network infrastructure overlaps with the fd97::<64 IPv6 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. The subnet for the distributed transit switch that enables east-west traffic. This subnet cannot overlap with any other subnets used by OVN-Kubernetes or on the host itself. It must be large enough to accommodate one IP address per node in your cluster.</p> <p>The default value is fd97::<64.</p>
internalJoinSubnet	string	<p>If your existing network infrastructure overlaps with the fd98::<64 IPv6 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. You must ensure that the IP address range does not overlap with any other subnet used by your OpenShift Container Platform installation. The IP address range must be larger than the maximum number of nodes that can be added to the cluster.</p> <p>The default value is fd98::<64.</p>

Table 18.52. `policyAuditConfig` object

Field	Type	Description
rateLimit	integer	The maximum number of messages to generate every second per node. The default value is 20 messages per second.
maxFileSize	integer	The maximum size for the audit log in bytes. The default value is 50000000 or 50 MB.
maxLogFiles	integer	The maximum number of log files that are retained.
destination	string	One of the following additional audit log targets: libc The libc syslog() function of the journald process on the host. udp:<host>:<port> A syslog server. Replace <host>:<port> with the host and port of the syslog server. unix:<file> A Unix Domain Socket file specified by <file> . null Do not send the audit logs to any additional target.
syslogFacility	string	The syslog facility, such as kern , as defined by RFC5424. The default value is local0 .

Table 18.53. gatewayConfig object

Field	Type	Description
routingViaHost	boolean	Set this field to true to send egress traffic from pods to the host networking stack. For highly-specialized installations and applications that rely on manually configured routes in the kernel routing table, you might want to route egress traffic to the host networking stack. By default, egress traffic is processed in OVN to exit the cluster and is not affected by specialized routes in the kernel routing table. The default value is false . This field has an interaction with the Open vSwitch hardware offloading feature. If you set this field to true , you do not receive the performance benefits of the offloading because egress traffic is processed by the host networking stack.
ipForwarding	object	You can control IP forwarding for all traffic on OVN-Kubernetes managed interfaces by using the ipForwarding specification in the Network resource. Specify Restricted to only allow IP forwarding for Kubernetes related traffic. Specify Global to allow forwarding of all IP traffic. For new installations, the default is Restricted . For updates to OpenShift Container Platform 4.14 or later, the default is Global .

Field	Type	Description
ipv4	object	Optional: Specify an object to configure the internal OVN-Kubernetes masquerade address for host to service traffic for IPv4 addresses.
ipv6	object	Optional: Specify an object to configure the internal OVN-Kubernetes masquerade address for host to service traffic for IPv6 addresses.

Table 18.54. gatewayConfig.ipv4 object

Field	Type	Description
internalMasqueradeSubnet	string	The masquerade IPv4 addresses that are used internally to enable host to service traffic. The host is configured with these IP addresses as well as the shared gateway bridge interface. The default value is 169.254.169.0/29 .

Table 18.55. gatewayConfig.ipv6 object

Field	Type	Description
internalMasqueradeSubnet	string	The masquerade IPv6 addresses that are used internally to enable host to service traffic. The host is configured with these IP addresses as well as the shared gateway bridge interface. The default value is fd69::/125 .

Table 18.56. ipsecConfig object

Field	Type	Description
mode	string	Specifies the behavior of the IPsec implementation. Must be one of the following values: <ul style="list-style-type: none"> ● Disabled: IPsec is not enabled on cluster nodes. ● External: IPsec is enabled for network traffic with external hosts. ● Full: IPsec is enabled for pod traffic and network traffic with external hosts.

Example OVN-Kubernetes configuration with IPsec enabled

```
defaultNetwork:
  type: OVNKubernetes
  ovnKubernetesConfig:
    mtu: 1400
```

```
genevePort: 6081
ipsecConfig:
mode: Full
```


**IMPORTANT**

Using OVNKubernetes can lead to a stack exhaustion problem on IBM Power®.

kubeProxyConfig object configuration (OpenShiftSDN container network interface only)

The values for the **kubeProxyConfig** object are defined in the following table:

Table 18.57. kubeProxyConfig object

Field	Type	Description
iptablesSyncPeriod	string	<p>The refresh period for iptables rules. The default value is 30s. Valid suffixes include s, m, and h and are described in the Go time package documentation.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p>NOTE</p> <p>Because of performance improvements introduced in OpenShift Container Platform 4.3 and greater, adjusting the iptablesSyncPeriod parameter is no longer necessary.</p> </div> </div>
proxyArguments.iptables-min-sync-period	array	<p>The minimum duration before refreshing iptables rules. This field ensures that the refresh does not happen too frequently. Valid suffixes include s, m, and h and are described in the Go time package. The default value is:</p> <pre>kubeProxyConfig: proxyArguments: iptables-min-sync-period: - 0s</pre>

18.4.11. Creating the Kubernetes manifest and Ignition config files

Because you must modify some cluster definition files and manually start the cluster machines, you must generate the Kubernetes manifest and Ignition config files that the cluster needs to configure the machines.

The installation configuration file transforms into the Kubernetes manifests. The manifests wrap into the Ignition configuration files, which are later used to configure the cluster machines.



IMPORTANT

- The Ignition config files that the OpenShift Container Platform installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.



NOTE

The installation program that generates the manifest and Ignition files is architecture specific and can be obtained from the [client image mirror](#). The Linux version of the installation program runs on s390x only. This installer program is also available as a Mac OS version.

Prerequisites

- You obtained the OpenShift Container Platform installation program.
- You created the **install-config.yaml** installation configuration file.

Procedure

1. Change to the directory that contains the OpenShift Container Platform installation program and generate the Kubernetes manifests for the cluster:

```
$ ./openshift-install create manifests --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the installation directory that contains the **install-config.yaml** file you created.



WARNING

If you are installing a three-node cluster, skip the following step to allow the control plane nodes to be schedulable.



IMPORTANT

When you configure control plane nodes from the default unschedulable to schedulable, additional subscriptions are required. This is because control plane nodes then become compute nodes.

2. Check that the **mastersSchedulable** parameter in the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** Kubernetes manifest file is set to **false**. This setting prevents pods from being scheduled on the control plane machines:
 - a. Open the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** file.
 - b. Locate the **mastersSchedulable** parameter and ensure that it is set to **false**.
 - c. Save and exit the file.
3. To create the Ignition configuration files, run the following command from the directory that contains the installation program:

```
$ ./openshift-install create ignition-configs --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the same installation directory.

Ignition config files are created for the bootstrap, control plane, and compute nodes in the installation directory. The **kubeadmin-password** and **kubeconfig** files are created in the **./<installation_directory>/auth** directory:

```
.
├── auth
│   ├── kubeadmin-password
│   └── kubeconfig
├── bootstrap.ign
├── master.ign
├── metadata.json
└── worker.ign
```

18.4.12. Installing RHCOS and starting the OpenShift Container Platform bootstrap process

To install OpenShift Container Platform on IBM Z[®] infrastructure that you provision, you must install Red Hat Enterprise Linux CoreOS (RHCOS) as Red Hat Enterprise Linux (RHEL) guest virtual machines. When you install RHCOS, you must provide the Ignition config file that was generated by the OpenShift Container Platform installation program for the type of machine you are installing. If you have configured suitable networking, DNS, and load balancing infrastructure, the OpenShift Container Platform bootstrap process begins automatically after the RHCOS machines have rebooted.

You can perform a fast-track installation of RHCOS that uses a prepackaged QEMU copy-on-write (QCOW2) disk image. Alternatively, you can perform a full installation on a new QCOW2 disk image.

To add further security to your system, you can optionally install RHCOS using IBM[®] Secure Execution before proceeding to the fast-track installation.

18.4.12.1. Installing RHCOS using IBM Secure Execution

Before you install RHCOS using IBM[®] Secure Execution, you must prepare the underlying infrastructure.

Prerequisites

- IBM[®] z15 or later, or IBM[®] LinuxONE III or later.

- Red Hat Enterprise Linux (RHEL) 8 or later.
- You have a bootstrap Ignition file. The file is not protected, enabling others to view and edit it.
- You have verified that the boot image has not been altered after installation.
- You must run all your nodes as IBM® Secure Execution guests.

Procedure

1. Prepare your RHEL KVM host to support IBM® Secure Execution.

- By default, KVM hosts do not support guests in IBM® Secure Execution mode. To support guests in IBM® Secure Execution mode, KVM hosts must boot in LPAR mode with the kernel parameter specification **prot_virt=1**. To enable **prot_virt=1** on RHEL 8, follow these steps:
 - a. Navigate to **/boot/loader/entries/** to modify your bootloader configuration file ***.conf**.
 - b. Add the kernel command line parameter **prot_virt=1**.
 - c. Run the **zipl** command and reboot your system.
KVM hosts that successfully start with support for IBM® Secure Execution for Linux issue the following kernel message:

```
prot_virt: Reserving <amount>MB as ultravisor base storage.
```

- d. To verify that the KVM host now supports IBM® Secure Execution, run the following command:

```
# cat /sys/firmware/uv/prot_virt_host
```

Example output

```
1
```

The value of this attribute is 1 for Linux instances that detect their environment as consistent with that of a secure host. For other instances, the value is 0.

2. Add your host keys to the KVM guest via Ignition.

During the first boot, RHCOS looks for your host keys to re-encrypt itself with them. RHCOS searches for files starting with **ibm-z-hostkey-** in the **/etc/se-hostkeys** directory. All host keys, for each machine the cluster is running on, must be loaded into the directory by the administrator. After first boot, you cannot run the VM on any other machines.



NOTE

You need to prepare your Ignition file on a safe system. For example, another IBM® Secure Execution guest.

For example:

```
{
  "ignition": { "version": "3.0.0" },
```

```

"storage": {
  "files": [
    {
      "path": "/etc/se-hostkeys/ibm-z-hostkey-<your-hostkey>.crt",
      "contents": {
        "source": "data:;base64,<base64 encoded hostkey document>"
      },
      "mode": 420
    },
    {
      "path": "/etc/se-hostkeys/ibm-z-hostkey-<your-hostkey>.crt",
      "contents": {
        "source": "data:;base64,<base64 encoded hostkey document>"
      },
      "mode": 420
    }
  ]
}
...

```



NOTE

You can add as many host keys as required if you want your node to be able to run on multiple IBM Z[®] machines.

- To generate the Base64 encoded string, run the following command:

```
base64 <your-hostkey>.crt
```

Compared to guests not running IBM[®] Secure Execution, the first boot of the machine is longer because the entire image is encrypted with a randomly generated LUKS passphrase before the Ignition phase.

- Add Ignition protection

To protect the secrets that are stored in the Ignition config file from being read or even modified, you must encrypt the Ignition config file.



NOTE

To achieve the desired security, Ignition logging and local login are disabled by default when running IBM[®] Secure Execution.

- Fetch the public GPG key for the **secex-qemu.qcow2** image and encrypt the Ignition config with the key by running the following command:

```
gpg --recipient-file /path/to/ignition.gpg.pub --yes --output /path/to/config.ign.gpg --verbose --armor --encrypt /path/to/config.ign
```

- Follow the fast-track installation of RHCOS to install nodes by using the IBM[®] Secure Execution QCOW image.

**NOTE**

Before you start the VM, replace **serial=ignition** with **serial=ignition_crypted**, and add the **launchSecurity** parameter.

Verification

When you have completed the fast-track installation of RHCOS and Ignition runs at the first boot, verify if decryption is successful.

- If the decryption is successful, you can expect an output similar to the following example:

Example output

```
[ 2.801433] systemd[1]: Starting coreos-ignition-setup-user.service - CoreOS Ignition User
Config Setup...

[ 2.803959] coreos-secex-ignition-decrypt[731]: gpg: key <key_name>: public key "Secure
Execution (secex) 38.20230323.dev.0" imported
[ 2.808874] coreos-secex-ignition-decrypt[740]: gpg: encrypted with rsa4096 key, ID
<key_name>, created <yyyy-mm-dd>
[ OK ] Finished coreos-secex-igni...S Secex Ignition Config Decryptor.
```

- If the decryption fails, you can expect an output similar to the following example:

Example output

```
Starting coreos-ignition-s...reOS Ignition User Config Setup...
[ 2.863675] coreos-secex-ignition-decrypt[729]: gpg: key <key_name>: public key "Secure
Execution (secex) 38.20230323.dev.0" imported
[ 2.869178] coreos-secex-ignition-decrypt[738]: gpg: encrypted with RSA key, ID
<key_name>
[ 2.870347] coreos-secex-ignition-decrypt[738]: gpg: public key decryption failed: No secret
key
[ 2.870371] coreos-secex-ignition-decrypt[738]: gpg: decryption failed: No secret key
```

Additional resources

- [Introducing IBM® Secure Execution for Linux](#)
- [Linux as an IBM® Secure Execution host or guest](#)
- [Setting up IBM® Secure Execution on IBM Z](#)

18.4.12.2. Configuring NBDE with static IP in an IBM Z or IBM LinuxONE environment

Enabling NBDE disk encryption in an IBM Z® or IBM® LinuxONE environment requires additional steps, which are described in detail in this section.

Prerequisites

- You have set up the External Tang Server. See [Network-bound disk encryption](#) for instructions.
- You have installed the **butane** utility.

- You have reviewed the instructions for how to create machine configs with Butane.

Procedure

- Create Butane configuration files for the control plane and compute nodes.
The following example of a Butane configuration for a control plane node creates a file named **master-storage.bu** for disk encryption:

```
variant: openshift
version: 4.16.0
metadata:
  name: master-storage
  labels:
    machineconfiguration.openshift.io/role: master
storage:
  luks:
    - clevis:
      tang:
        - thumbprint: QcPr_NHFJammnRCA3fFMVdNBwjs
          url: http://clevis.example.com:7500
      options: ❶
        - --cipher
          - aes-cbc-essiv:sha256
    device: /dev/disk/by-partlabel/root
    label: luks-root
    name: root
    wipe_volume: true
  filesystems:
    - device: /dev/mapper/root
      format: xfs
      label: root
      wipe_filesystem: true
openshift:
  fips: true ❷
```

- The cipher option is only required if FIPS mode is enabled. Omit the entry if FIPS is disabled.
- Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.

- Create a customized initramfs file to boot the machine, by running the following command:

```
$ coreos-installer pxe customize \
  /root/rhcos-bootfiles/rhcos-<release>-live-initramfs.s390x.img \
  --dest-device /dev/disk/by-id/scsi-<serial_number> --dest-karg-append \
  ip=<ip_address>::<gateway_ip>:<subnet_mask>::<network_device>:none \
  --dest-karg-append nameserver=<nameserver_ip> \
  --dest-karg-append rd.neednet=1 -o \
  /root/rhcos-bootfiles/<node_name>-initramfs.s390x.img
```

**NOTE**

Before first boot, you must customize the `initramfs` for each node in the cluster, and add PXE kernel parameters.

3. Create a parameter file that includes **`ignition.platform.id=metal`** and **`ignition.firstboot`**.

```
rd.neednet=1 \
console=ttySCLP0 \
ignition.firstboot ignition.platform.id=metal \
coreos.live.rootfs_url=http://<http_server>/rhcos-<version>-live-rootfs.<architecture>.img \ 1
coreos.inst.ignition_url=http://<http_server>/master.ign \ 2
ip=10.19.17.2::10.19.17.1:255.255.255.0::enb0d0:none nameserver=10.19.17.1 \
zfcplib.allow_lun_scan=0 \
rd.znet=qeth,0.0.bdd0,0.0.bdd1,0.0.bdd2,layer2=1 \
rd.zfcp=0.0.5677,0x600606680g7f0056,0x034F000000000000
```

- 1 Specify the location of the **rootfs** artifact for the **kernel** and **initramfs** you are booting. Only HTTP and HTTPS protocols are supported.
- 2 Specify the location of the Ignition config file. Use **master.ign** or **worker.ign**. Only HTTP and HTTPS protocols are supported.

**NOTE**

Write all options in the parameter file as a single line and make sure you have no newline characters.

Additional resources

- [Creating machine configs with Butane](#)

18.4.12.3. Fast-track installation by using a prepackaged QCOW2 disk image


Complete the following steps to create the machines in a fast-track installation of Red Hat Enterprise Linux CoreOS (RHCOS), importing a prepackaged Red Hat Enterprise Linux CoreOS (RHCOS) QEMU copy-on-write (QCOW2) disk image.

Prerequisites

- At least one LPAR running on RHEL 8.6 or later with KVM, referred to as RHEL KVM host in this procedure.
- The KVM/QEMU hypervisor is installed on the RHEL KVM host.
- A domain name server (DNS) that can perform hostname and reverse lookup for the nodes.
- A DHCP server that provides IP addresses.

Procedure

1. Obtain the RHEL QEMU copy-on-write (QCOW2) disk image file from the [Product Downloads](#) page on the Red Hat Customer Portal or from the [RHCOS image mirror](#) page.

**IMPORTANT**

The RHCOS images might not change with every release of OpenShift Container Platform. You must download images with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Only use the appropriate RHCOS QCOW2 image described in the following procedure.

- Download the QCOW2 disk image and Ignition files to a common directory on the RHEL KVM host.

For example: **`/var/lib/libvirt/images`**

**NOTE**

The Ignition files are generated by the OpenShift Container Platform installer.

- Create a new disk image with the QCOW2 disk image backing file for each KVM guest node.

```
$ qemu-img create -f qcow2 -F qcow2 -b /var/lib/libvirt/images/{source_rhcos_qemu}
/var/lib/libvirt/images/{vmname}.qcow2 {size}
```

- Create the new KVM guest nodes using the Ignition file and the new disk image.

```
$ virt-install --noautoconsole \
  --connect qemu:///system \
  --name {vm_name} \
  --memory {memory} \
  --vcpus {vcpus} \
  --disk {disk} \
  --launchSecurity type="s390-pv" 1 \
  --import \
  --network network={network},mac={mac} \
  --disk path={ign_file},format=raw,readonly=on,serial=ignition,startup_policy=optional 2
```

- 1** If IBM® Secure Execution is enabled, add the **`launchSecurity type="s390-pv"`** parameter.
- 2** If IBM® Secure Execution is enabled, replace **`serial=ignition`** with **`serial=ignition_crypted`**.

18.4.12.4. Full installation on a new QCOW2 disk image

Complete the following steps to create the machines in a full installation on a new QEMU copy-on-write (QCOW2) disk image.

Prerequisites

- At least one LPAR running on RHEL 8.6 or later with KVM, referred to as RHEL KVM host in this procedure.
- The KVM/QEMU hypervisor is installed on the RHEL KVM host.
- A domain name server (DNS) that can perform hostname and reverse lookup for the nodes.
- An HTTP or HTTPS server is set up.

Procedure

1. Obtain the RHEL kernel, initramfs, and rootfs files from the [Product Downloads](#) page on the Red Hat Customer Portal or from the [RHCOS image mirror](#) page.



IMPORTANT

The RHCOS images might not change with every release of OpenShift Container Platform. You must download images with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Only use the appropriate RHCOS QCOW2 image described in the following procedure.

The file names contain the OpenShift Container Platform version number. They resemble the following examples:

- kernel: **rhcos-`<version>`-live-kernel-`<architecture>`**
 - initramfs: **rhcos-`<version>`-live-initramfs.`<architecture>`.img**
 - rootfs: **rhcos-`<version>`-live-rootfs.`<architecture>`.img**
2. Move the downloaded RHEL live kernel, initramfs, and rootfs as well as the Ignition files to an HTTP or HTTPS server before you launch **virt-install**.



NOTE

The Ignition files are generated by the OpenShift Container Platform installer.

3. Create the new KVM guest nodes using the RHEL kernel, initramfs, and Ignition files, the new disk image, and adjusted parm line arguments.
 - For **--location**, specify the location of the kernel/initrd on the HTTP or HTTPS server.
 - For **coreos.inst.ignition_url=**, specify the Ignition file for the machine role. Use **bootstrap.ign**, **master.ign**, or **worker.ign**. Only HTTP and HTTPS protocols are supported.
 - For **coreos.live.rootfs_url=**, specify the matching rootfs artifact for the kernel and initramfs you are booting. Only HTTP and HTTPS protocols are supported.

```
$ virt-install \
  --connect qemu:///system \
  --name {vm_name} \
  --vcpus {vcpus} \
  --memory {memory_mb} \
  --disk {vm_name}.qcow2,size={image_size| default(10,true)} \
  --network network={virt_network_parm} \
  --boot hd \
  --location {media_location},kernel={rhcos_kernel},initrd={rhcos_initrd} \
  --extra-args "rd.neednet=1 coreos.inst.install_dev=/dev/vda coreos.live.rootfs_url=
{rhcos_liveos} ip={ip}::{default_gateway}:{subnet_mask_length}:{vm_name}:enc1:none:
{MTU} nameserver={dns} coreos.inst.ignition_url={rhcos_ign}" \
  --noautoconsole \
  --wait
```

18.4.12.5. Advanced RHCOS installation reference

This section illustrates the networking configuration and other advanced options that allow you to modify the Red Hat Enterprise Linux CoreOS (RHCOS) manual installation process. The following tables describe the kernel arguments and command-line options you can use with the RHCOS live installer and the **coreos-installer** command.

18.4.12.5.1. Networking options for ISO installations

If you install RHCOS from an ISO image, you can add kernel arguments manually when you boot the image to configure networking for a node. If no networking arguments are specified, DHCP is activated in the initramfs when RHCOS detects that networking is required to fetch the Ignition config file.



IMPORTANT

When adding networking arguments manually, you must also add the **rd.neednet=1** kernel argument to bring the network up in the initramfs.

The following information provides examples for configuring networking on your RHCOS nodes for ISO installations. The examples describe how to use the **ip=** and **nameserver=** kernel arguments.



NOTE

Ordering is important when adding the kernel arguments: **ip=** and **nameserver=**.

The networking options are passed to the **dracut** tool during system boot. For more information about the networking options supported by **dracut**, see the **dracut.cmdline** manual page.

The following examples are the networking options for ISO installation.

Configuring DHCP or static IP addresses

To configure an IP address, either use DHCP (**ip=dhcp**) or set an individual static IP address (**ip=<host_ip>**). If setting a static IP, you must then identify the DNS server IP address (**nameserver=<dns_ip>**) on each node. The following example sets:

- The node's IP address to **10.10.10.2**
- The gateway address to **10.10.10.254**
- The netmask to **255.255.255.0**
- The hostname to **core0.example.com**
- The DNS server address to **4.4.4.41**
- The auto-configuration value to **none**. No auto-configuration is required when IP networking is configured statically.

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none  
nameserver=4.4.4.41
```



NOTE

When you use DHCP to configure IP addressing for the RHCOS machines, the machines also obtain the DNS server information through DHCP. For DHCP-based deployments, you can define the DNS server address that is used by the RHCOS nodes through your DHCP server configuration.

Configuring an IP address without a static hostname

You can configure an IP address without assigning a static hostname. If a static hostname is not set by the user, it will be picked up and automatically set by a reverse DNS lookup. To configure an IP address without a static hostname refer to the following example:

- The node's IP address to **10.10.10.2**
- The gateway address to **10.10.10.254**
- The netmask to **255.255.255.0**
- The DNS server address to **4.4.4.41**
- The auto-configuration value to **none**. No auto-configuration is required when IP networking is configured statically.

```
ip=10.10.10.2::10.10.10.254:255.255.255.0::enp1s0:none
nameserver=4.4.4.41
```

Specifying multiple network interfaces

You can specify multiple network interfaces by setting multiple **ip=** entries.

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none
ip=10.10.10.3::10.10.10.254:255.255.255.0:core0.example.com:enp2s0:none
```

Configuring default gateway and route

Optional: You can configure routes to additional networks by setting an **rd.route=** value.



NOTE

When you configure one or multiple networks, one default gateway is required. If the additional network gateway is different from the primary network gateway, the default gateway must be the primary network gateway.

- Run the following command to configure the default gateway:

```
ip=::10.10.10.254:::
```

- Enter the following command to configure the route for the additional network:

```
rd.route=20.20.20.0/24:20.20.20.254:enp2s0
```

Disabling DHCP on a single interface

You can disable DHCP on a single interface, such as when there are two or more network interfaces and only one interface is being used. In the example, the **enp1s0** interface has a static networking configuration and DHCP is disabled for **enp2s0**, which is not used:

■

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none  
ip=:::core0.example.com:enp2s0:none
```

Combining DHCP and static IP configurations

You can combine DHCP and static IP configurations on systems with multiple network interfaces, for example:

```
ip=enp1s0:dhcp  
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp2s0:none
```

Configuring VLANs on individual interfaces

Optional: You can configure VLANs on individual interfaces by using the **vlan=** parameter.

- To configure a VLAN on a network interface and use a static IP address, run the following command:

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp2s0.100:none  
vlan=enp2s0.100:enp2s0
```

- To configure a VLAN on a network interface and to use DHCP, run the following command:

```
ip=enp2s0.100:dhcp  
vlan=enp2s0.100:enp2s0
```

Providing multiple DNS servers

You can provide multiple DNS servers by adding a **nameserver=** entry for each server, for example:

```
nameserver=1.1.1.1  
nameserver=8.8.8.8
```

18.4.13. Waiting for the bootstrap process to complete

The OpenShift Container Platform bootstrap process begins after the cluster nodes first boot into the persistent RHCOS environment that has been installed to disk. The configuration information provided through the Ignition config files is used to initialize the bootstrap process and install OpenShift Container Platform on the machines. You must wait for the bootstrap process to complete.

Prerequisites

- You have created the Ignition config files for your cluster.
- You have configured suitable network, DNS and load balancing infrastructure.
- You have obtained the installation program and generated the Ignition config files for your cluster.
- You installed RHCOS on your cluster machines and provided the Ignition config files that the OpenShift Container Platform installation program generated.
- Your machines have direct internet access or have an HTTP or HTTPS proxy available.

Procedure

1. Monitor the bootstrap process:


```
$ ./openshift-install --dir <installation_directory> wait-for bootstrap-complete \ 1
--log-level=info 2
```

1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Example output

```
INFO Waiting up to 30m0s for the Kubernetes API at https://api.test.example.com:6443...
INFO API v1.29.4 up
INFO Waiting up to 30m0s for bootstrapping to complete...
INFO It is now safe to remove the bootstrap resources
```

The command succeeds when the Kubernetes API server signals that it has been bootstrapped on the control plane machines.

2. After the bootstrap process is complete, remove the bootstrap machine from the load balancer.



IMPORTANT

You must remove the bootstrap machine from the load balancer at this point. You can also remove or reformat the bootstrap machine itself.

18.4.14. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

18.4.15. Approving the certificate signing requests for your machines

When you add machines to a cluster, two pending certificate signing requests (CSRs) are generated for each machine that you added. You must confirm that these CSRs are approved or, if necessary, approve them yourself. The client requests must be approved first, followed by the server requests.

Prerequisites

- You added machines to your cluster.

Procedure

- Confirm that the cluster recognizes the machines:

```
$ oc get nodes
```

Example output

NAME	STATUS	ROLES	AGE	VERSION
master-0	Ready	master	63m	v1.29.4
master-1	Ready	master	63m	v1.29.4
master-2	Ready	master	64m	v1.29.4

The output lists all of the machines that you created.



NOTE

The preceding output might not include the compute nodes, also known as worker nodes, until some CSRs are approved.

- Review the pending CSRs and ensure that you see the client requests with the **Pending** or **Approved** status for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

NAME	AGE	REQUESTOR	CONDITION
csr-mddf5	20m	system:node:master-01.example.com	Approved,Issued
csr-z5rln	16m	system:node:worker-21.example.com	Approved,Issued

- If the CSRs were not approved, after all of the pending CSRs for the machines you added are in **Pending** status, approve the CSRs for your cluster machines:

**NOTE**

Because the CSRs rotate automatically, approve your CSRs within an hour of adding the machines to the cluster. If you do not approve them within an hour, the certificates will rotate, and more than two certificates will be present for each node. You must approve all of these certificates. After the client CSR is approved, the Kubelet creates a secondary CSR for the serving certificate, which requires manual approval. Then, subsequent serving certificate renewal requests are automatically approved by the **machine-approver** if the Kubelet requests a new certificate with identical parameters.

**NOTE**

For clusters running on platforms that are not machine API enabled, such as bare metal and other user-provisioned infrastructure, you must implement a method of automatically approving the kubelet serving certificate requests (CSRs). If a request is not approved, then the **oc exec**, **oc rsh**, and **oc logs** commands cannot succeed, because a serving certificate is required when the API server connects to the kubelet. Any operation that contacts the Kubelet endpoint requires this certificate approval to be in place. The method must watch for new CSRs, confirm that the CSR was submitted by the **node-bootstrap** service account in the **system:node** or **system:admin** groups, and confirm the identity of the node.

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{"\n"}{{end}}{{end}}' | xargs --no-run-if-empty oc adm certificate approve
```

**NOTE**

Some Operators might not become available until some CSRs are approved.

4. Now that your client requests are approved, you must review the server requests for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

```
NAME          AGE   REQUESTOR                                     CONDITION
csr-bfd72     5m26s system:node:ip-10-0-50-126.us-east-2.compute.internal
Pending
csr-c57lv     5m26s system:node:ip-10-0-95-157.us-east-2.compute.internal
Pending
...
```

- If the remaining CSRs are not approved, and are in the **Pending** status, approve the CSRs for your cluster machines:

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** `<csr_name>` is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{\n"}\n{{end}}' | xargs oc adm certificate approve
```

- After all client and server CSRs have been approved, the machines have the **Ready** status. Verify this by running the following command:

```
$ oc get nodes
```

Example output

```
NAME      STATUS  ROLES  AGE  VERSION
master-0  Ready   master 73m  v1.29.4
master-1  Ready   master 73m  v1.29.4
master-2  Ready   master 74m  v1.29.4
worker-0  Ready   worker 11m  v1.29.4
worker-1  Ready   worker 11m  v1.29.4
```



NOTE

It can take a few minutes after approval of the server CSRs for the machines to transition to the **Ready** status.

Additional information

- For more information on CSRs, see [Certificate Signing Requests](#).

18.4.16. Initial Operator configuration

After the control plane initializes, you must immediately configure some Operators so that they all become available.

Prerequisites

- Your control plane has initialized.

Procedure

- Watch the cluster components come online:

```
$ watch -n5 oc get clusteroperators
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.16.0	True	False	False	19m
baremetal	4.16.0	True	False	False	37m
cloud-credential	4.16.0	True	False	False	40m
cluster-autoscaler	4.16.0	True	False	False	37m
config-operator	4.16.0	True	False	False	38m
console	4.16.0	True	False	False	26m
csi-snapshot-controller	4.16.0	True	False	False	37m
dns	4.16.0	True	False	False	37m
etcd	4.16.0	True	False	False	36m
image-registry	4.16.0	True	False	False	31m
ingress	4.16.0	True	False	False	30m
insights	4.16.0	True	False	False	31m
kube-apiserver	4.16.0	True	False	False	26m
kube-controller-manager	4.16.0	True	False	False	36m
kube-scheduler	4.16.0	True	False	False	36m
kube-storage-version-migrator	4.16.0	True	False	False	37m
machine-api	4.16.0	True	False	False	29m
machine-approver	4.16.0	True	False	False	37m
machine-config	4.16.0	True	False	False	36m
marketplace	4.16.0	True	False	False	37m
monitoring	4.16.0	True	False	False	29m
network	4.16.0	True	False	False	38m
node-tuning	4.16.0	True	False	False	37m
openshift-apiserver	4.16.0	True	False	False	32m
openshift-controller-manager	4.16.0	True	False	False	30m
openshift-samples	4.16.0	True	False	False	32m
operator-lifecycle-manager	4.16.0	True	False	False	37m
operator-lifecycle-manager-catalog	4.16.0	True	False	False	37m
operator-lifecycle-manager-packageserver	4.16.0	True	False	False	32m
service-ca	4.16.0	True	False	False	38m
storage	4.16.0	True	False	False	37m

2. Configure the Operators that are not available.

18.4.16.1. Image registry storage configuration

The Image Registry Operator is not initially available for platforms that do not provide default storage. After installation, you must configure your registry to use storage so that the Registry Operator is made available.

Instructions are shown for configuring a persistent volume, which is required for production clusters. Where applicable, instructions are shown for configuring an empty directory as the storage location, which is available for only non-production clusters.


Additional instructions are provided for allowing the image registry to use block storage types by using the **Recreate** rollout strategy during upgrades.

18.4.16.1.1. Configuring registry storage for IBM Z

As a cluster administrator, following installation you must configure your registry to use storage.

Prerequisites

- You have access to the cluster as a user with the **cluster-admin** role.
- You have a cluster on IBM Z®.
- You have provisioned persistent storage for your cluster, such as Red Hat OpenShift Data Foundation.



IMPORTANT

OpenShift Container Platform supports **ReadWriteOnce** access for image registry storage when you have only one replica. **ReadWriteOnce** access also requires that the registry uses the **Recreate** rollout strategy. To deploy an image registry that supports high availability with two or more replicas, **ReadWriteMany** access is required.

- Must have 100Gi capacity.

Procedure

1. To configure your registry to use storage, change the **spec.storage.pvc** in the **configs.imageregistry/cluster** resource.



NOTE

When you use shared storage, review your security settings to prevent outside access.

2. Verify that you do not have a registry pod:

```
$ oc get pod -n openshift-image-registry -l docker-registry=default
```

Example output

```
No resources found in openshift-image-registry namespace
```



NOTE

If you do have a registry pod in your output, you do not need to continue with this procedure.

3. Check the registry configuration:

```
$ oc edit configs.imageregistry.operator.openshift.io
```

Example output

```
storage:
  pvc:
    claim:
```

Leave the **claim** field blank to allow the automatic creation of an **image-registry-storage** PVC.

4. Check the **clusteroperator** status:

```
$ oc get clusteroperator image-registry
```

Example output

```
NAME          VERSION          AVAILABLE PROGRESSING DEGRADED SINCE
MESSAGE
image-registry 4.16             True      False      False  6h50m
```

5. Ensure that your registry is set to managed to enable building and pushing of images.

- Run:

```
$ oc edit configs.imageregistry/cluster
```

Then, change the line

```
managementState: Removed
```

to

```
managementState: Managed
```

18.4.16.1.2. Configuring storage for the image registry in non-production clusters

You must configure storage for the Image Registry Operator. For non-production clusters, you can set the image registry to an empty directory. If you do so, all images are lost if you restart the registry.

Procedure

- To set the image registry storage to an empty directory:

```
$ oc patch configs.imageregistry.operator.openshift.io cluster --type merge --patch '{"spec": {"storage":{"emptyDir":{}}}'
```



WARNING

Configure this option for only non-production clusters.

If you run this command before the Image Registry Operator initializes its components, the **oc patch** command fails with the following error:

```
Error from server (NotFound): configs.imageregistry.operator.openshift.io "cluster" not found
```

Wait a few minutes and run the command again.

18.4.17. Completing installation on user-provisioned infrastructure

After you complete the Operator configuration, you can finish installing the cluster on infrastructure that you provide.

Prerequisites

- Your control plane has initialized.
- You have completed the initial Operator configuration.

Procedure

1. Confirm that all the cluster components are online with the following command:

```
$ watch -n5 oc get clusteroperators
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.16.0	True	False	False	19m
baremetal	4.16.0	True	False	False	37m
cloud-credential	4.16.0	True	False	False	40m
cluster-autoscaler	4.16.0	True	False	False	37m
config-operator	4.16.0	True	False	False	38m
console	4.16.0	True	False	False	26m
csi-snapshot-controller	4.16.0	True	False	False	37m
dns	4.16.0	True	False	False	37m
etcd	4.16.0	True	False	False	36m
image-registry	4.16.0	True	False	False	31m
ingress	4.16.0	True	False	False	30m
insights	4.16.0	True	False	False	31m
kube-apiserver	4.16.0	True	False	False	26m
kube-controller-manager	4.16.0	True	False	False	36m
kube-scheduler	4.16.0	True	False	False	36m
kube-storage-version-migrator	4.16.0	True	False	False	37m
machine-api	4.16.0	True	False	False	29m
machine-approver	4.16.0	True	False	False	37m
machine-config	4.16.0	True	False	False	36m
marketplace	4.16.0	True	False	False	37m
monitoring	4.16.0	True	False	False	29m
network	4.16.0	True	False	False	38m
node-tuning	4.16.0	True	False	False	37m
openshift-apiserver	4.16.0	True	False	False	32m
openshift-controller-manager	4.16.0	True	False	False	30m
openshift-samples	4.16.0	True	False	False	32m
operator-lifecycle-manager	4.16.0	True	False	False	37m
operator-lifecycle-manager-catalog	4.16.0	True	False	False	37m
operator-lifecycle-manager-packageserver	4.16.0	True	False	False	32m
service-ca	4.16.0	True	False	False	38m
storage	4.16.0	True	False	False	37m

Alternatively, the following command notifies you when all of the clusters are available. It also retrieves and displays credentials:

```
$ ./openshift-install --dir <installation_directory> wait-for install-complete 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

Example output

```
INFO Waiting up to 30m0s for the cluster to initialize...
```

The command succeeds when the Cluster Version Operator finishes deploying the OpenShift Container Platform cluster from Kubernetes API server.



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

2. Confirm that the Kubernetes API server is communicating with the pods.

- a. To view a list of all pods, use the following command:

```
$ oc get pods --all-namespaces
```

Example output

```

NAMESPACE          NAME                                     READY  STATUS
RESTARTS  AGE
openshift-apiserver-operator  openshift-apiserver-operator-85cb746d55-zqhs8  1/1
Running    1    9m
openshift-apiserver          apiserver-67b9g                                1/1  Running  0
3m
openshift-apiserver          apiserver-ljcmx                                1/1  Running  0
1m
openshift-apiserver          apiserver-z25h4                                1/1  Running  0
2m
openshift-authentication-operator  authentication-operator-69d5d8bf84-vh2n8    1/1
Running    0    5m
...
```

-
- b. View the logs for a pod that is listed in the output of the previous command by using the following command:

```
$ oc logs <pod_name> -n <namespace> 1
```

- 1 Specify the pod name and namespace, as shown in the output of the previous command.

If the pod logs display, the Kubernetes API server can communicate with the cluster machines.

3. For an installation with Fibre Channel Protocol (FCP), additional steps are required to enable multipathing. Do not enable multipathing during installation. See "Enabling multipathing with kernel arguments on RHCOS" in the *Postinstallation machine configuration tasks* documentation for more information.

18.4.18. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources


- See [About remote health monitoring](#) for more information about the Telemetry service
- [How to generate SOSREPORT within OpenShift4 nodes without SSH](#) .

18.4.19. Next steps

- [Customize your cluster](#).
- If necessary, you can [opt out of remote health reporting](#) .

18.5. INSTALLING A CLUSTER WITH RHEL KVM ON IBM Z AND IBM LINUXONE IN A RESTRICTED NETWORK

In OpenShift Container Platform version 4.16, you can install a cluster on IBM Z® or IBM® LinuxONE infrastructure that you provision in a restricted network.



NOTE

While this document refers to only IBM Z®, all information in it also applies to IBM® LinuxONE.

**IMPORTANT**

Additional considerations exist for non-bare metal platforms. Review the information in the [guidelines for deploying OpenShift Container Platform on non-tested platforms](#) before you install an OpenShift Container Platform cluster.

18.5.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [created a registry on your mirror host](#) and obtained the **imageContentSources** data for your version of OpenShift Container Platform.
- You must move or remove any existing installation files, before you begin the installation process. This ensures that the required installation files are created and updated during the installation process.

**IMPORTANT**

Ensure that installation steps are done from a machine with access to the installation media.

- You provisioned [persistent storage using OpenShift Data Foundation](#) or other supported storage protocols for your cluster. To deploy a private image registry, you must set up persistent storage with **ReadWriteMany** access.
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.

**NOTE**

Be sure to also review this site list if you are configuring a proxy.

- You provisioned a RHEL Kernel Virtual Machine (KVM) system that is hosted on the logical partition (LPAR) and based on RHEL 8.6 or later. See [Red Hat Enterprise Linux 8 and 9 Life Cycle](#).

18.5.2. About installations in restricted networks

In OpenShift Container Platform 4.16, you can perform an installation that does not require an active connection to the internet to obtain software components. Restricted network installations can be completed using installer-provisioned infrastructure or user-provisioned infrastructure, depending on the cloud platform to which you are installing the cluster.

If you choose to perform a restricted network installation on a cloud platform, you still require access to its cloud APIs. Some cloud functions, like Amazon Web Service's Route 53 DNS and IAM services, require internet access. Depending on your network, you might require less internet access for an installation on bare metal hardware, Nutanix, or on VMware vSphere.

To complete a restricted network installation, you must create a registry that mirrors the contents of the OpenShift image registry and contains the installation media. You can create this registry on a mirror

host, which can access both the internet and your closed network, or by using other methods that meet your restrictions.



IMPORTANT

Because of the complexity of the configuration for user-provisioned installations, consider completing a standard user-provisioned infrastructure installation before you attempt a restricted network installation using user-provisioned infrastructure. Completing this test installation might make it easier to isolate and troubleshoot any issues that might arise during your installation in a restricted network.

18.5.2.1. Additional limits

Clusters in restricted networks have the following additional limitations and restrictions:

- The **ClusterVersion** status includes an **Unable to retrieve available updates** error.
- By default, you cannot use the contents of the Developer Catalog because you cannot access the required image stream tags.

18.5.3. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to obtain the images that are necessary to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.

18.5.4. Machine requirements for a cluster with user-provisioned infrastructure

For a cluster that contains user-provisioned infrastructure, you must deploy all of the required machines.

One or more KVM host machines based on RHEL 8.6 or later. Each RHEL KVM host machine must have libvirt installed and running. The virtual machines are provisioned under each RHEL KVM host machine.

18.5.4.1. Required machines

The smallest OpenShift Container Platform clusters require the following hosts:

Table 18.58. Minimum required hosts

Hosts	Description
One temporary bootstrap machine	The cluster requires the bootstrap machine to deploy the OpenShift Container Platform cluster on the three control plane machines. You can remove the bootstrap machine after you install the cluster.

Hosts	Description
Three control plane machines	The control plane machines run the Kubernetes and OpenShift Container Platform services that form the control plane.
At least two compute machines, which are also known as worker machines.	The workloads requested by OpenShift Container Platform users run on the compute machines.



IMPORTANT

To improve high availability of your cluster, distribute the control plane machines over different RHEL instances on at least two physical machines.

The bootstrap, control plane, and compute machines must use Red Hat Enterprise Linux CoreOS (RHCOS) as the operating system.

See [Red Hat Enterprise Linux technology capabilities and limits](#) .

18.5.4.2. Network connectivity requirements

The OpenShift Container Platform installer creates the Ignition files, which are necessary for all the Red Hat Enterprise Linux CoreOS (RHCOS) virtual machines. The automated installation of OpenShift Container Platform is performed by the bootstrap machine. It starts the installation of OpenShift Container Platform on each node, starts the Kubernetes cluster, and then finishes. During this bootstrap, the virtual machine must have an established network connection either through a Dynamic Host Configuration Protocol (DHCP) server or static IP address.

18.5.4.3. IBM Z network connectivity requirements

To install on IBM Z[®] under RHEL KVM, you need:

- A RHEL KVM host configured with an OSA or RoCE network adapter.
- Either a RHEL KVM host that is configured to use bridged networking in libvirt or MacVTap to connect the network to the guests.

See [Types of virtual network connections](#) .

18.5.4.4. Host machine resource requirements

The RHEL KVM host in your environment must meet the following requirements to host the virtual machines that you plan for the OpenShift Container Platform environment. See [Getting started with virtualization](#).

You can install OpenShift Container Platform version 4.16 on the following IBM[®] hardware:

- IBM[®] z16 (all models), IBM[®] z15 (all models), IBM[®] z14 (all models)
- IBM[®] LinuxONE 4 (all models), IBM[®] LinuxONE III (all models), IBM[®] LinuxONE Emperor II, IBM[®] LinuxONE Rockhopper II

18.5.4.5. Minimum IBM Z system environment

Hardware requirements

- The equivalent of six Integrated Facilities for Linux (IFL), which are SMT2 enabled, for each cluster.
- At least one network connection to both connect to the **LoadBalancer** service and to serve data for traffic outside the cluster.



NOTE

You can use dedicated or shared IFLs to assign sufficient compute resources. Resource sharing is one of the key strengths of IBM Z®. However, you must adjust capacity correctly on each hypervisor layer and ensure sufficient resources for every OpenShift Container Platform cluster.



IMPORTANT

Since the overall performance of the cluster can be impacted, the LPARs that are used to set up the OpenShift Container Platform clusters must provide sufficient compute capacity. In this context, LPAR weight management, entitlements, and CPU shares on the hypervisor level play an important role.

Operating system requirements

- One LPAR running on RHEL 8.6 or later with KVM, which is managed by libvirt

On your RHEL KVM host, set up:

- Three guest virtual machines for OpenShift Container Platform control plane machines
- Two guest virtual machines for OpenShift Container Platform compute machines
- One guest virtual machine for the temporary OpenShift Container Platform bootstrap machine

18.5.4.6. Minimum resource requirements

Each cluster virtual machine must meet the following minimum requirements:

Virtual Machine	Operating System	vCPU [1]	Virtual RAM	Storage	IOPS
Bootstrap	RHCOS	4	16 GB	100 GB	N/A
Control plane	RHCOS	4	16 GB	100 GB	N/A
Compute	RHCOS	2	8 GB	100 GB	N/A

1. One physical core (IFL) provides two logical cores (threads) when SMT-2 is enabled. The hypervisor can provide two or more vCPUs.

18.5.4.7. Preferred IBM Z system environment

Hardware requirements

- Three LPARS that each have the equivalent of six IFLs, which are SMT2 enabled, for each cluster.
- Two network connections to both connect to the **LoadBalancer** service and to serve data for traffic outside the cluster.

Operating system requirements

- For high availability, two or three LPARs running on RHEL 8.6 or later with KVM, which are managed by libvirt.

On your RHEL KVM host, set up:

- Three guest virtual machines for OpenShift Container Platform control plane machines, distributed across the RHEL KVM host machines.
- At least six guest virtual machines for OpenShift Container Platform compute machines, distributed across the RHEL KVM host machines.
- One guest virtual machine for the temporary OpenShift Container Platform bootstrap machine.
- To ensure the availability of integral components in an overcommitted environment, increase the priority of the control plane by using **cpu_shares**. Do the same for infrastructure nodes, if they exist. See [schedinfo](#) in IBM® Documentation.

18.5.4.8. Preferred resource requirements

The preferred requirements for each cluster virtual machine are:

Virtual Machine	Operating System	vCPU	Virtual RAM	Storage
Bootstrap	RHCOS	4	16 GB	120 GB
Control plane	RHCOS	8	16 GB	120 GB
Compute	RHCOS	6	8 GB	120 GB

18.5.4.9. Certificate signing requests management

Because your cluster has limited access to automatic machine management when you use infrastructure that you provision, you must provide a mechanism for approving cluster certificate signing requests (CSRs) after installation. The **kube-controller-manager** only approves the kubelet client CSRs. The **machine-approver** cannot guarantee the validity of a serving certificate that is requested by using kubelet credentials because it cannot confirm that the correct machine issued the request. You must determine and implement a method of verifying the validity of the kubelet serving certificate requests and approving them.

Additional resources

- [Recommended host practices for IBM Z® & IBM® LinuxONE environments](#)

18.5.4.10. Networking requirements for user-provisioned infrastructure

All the Red Hat Enterprise Linux CoreOS (RHCOS) machines require networking to be configured in **initramfs** during boot to fetch their Ignition config files.

During the initial boot, the machines require an IP address configuration that is set either through a DHCP server or statically by providing the required boot options. After a network connection is established, the machines download their Ignition config files from an HTTP or HTTPS server. The Ignition config files are then used to set the exact state of each machine. The Machine Config Operator completes more changes to the machines, such as the application of new certificates or keys, after installation.

It is recommended to use a DHCP server for long-term management of the cluster machines. Ensure that the DHCP server is configured to provide persistent IP addresses, DNS server information, and hostnames to the cluster machines.



NOTE

If a DHCP service is not available for your user-provisioned infrastructure, you can instead provide the IP networking configuration and the address of the DNS server to the nodes at RHCOS install time. These can be passed as boot arguments if you are installing from an ISO image. See the *Installing RHCOS and starting the OpenShift Container Platform bootstrap process* section for more information about static IP provisioning and advanced networking options.

The Kubernetes API server must be able to resolve the node names of the cluster machines. If the API servers and worker nodes are in different zones, you can configure a default DNS search zone to allow the API server to resolve the node names. Another supported approach is to always refer to hosts by their fully-qualified domain names in both the node objects and all DNS requests.

18.5.4.10.1. Setting the cluster node hostnames through DHCP

On Red Hat Enterprise Linux CoreOS (RHCOS) machines, the hostname is set through NetworkManager. By default, the machines obtain their hostname through DHCP. If the hostname is not provided by DHCP, set statically through kernel arguments, or another method, it is obtained through a reverse DNS lookup. Reverse DNS lookup occurs after the network has been initialized on a node and can take time to resolve. Other system services can start prior to this and detect the hostname as **localhost** or similar. You can avoid this by using DHCP to provide the hostname for each cluster node.

Additionally, setting the hostnames through DHCP can bypass any manual DNS record name configuration errors in environments that have a DNS split-horizon implementation.

18.5.4.10.2. Network connectivity requirements

You must configure the network connectivity between machines to allow OpenShift Container Platform cluster components to communicate. Each machine must be able to resolve the hostnames of all other machines in the cluster.

This section provides details about the ports that are required.

Table 18.59. Ports used for all-machine to all-machine communications

Protocol	Port	Description
ICMP	N/A	Network reachability tests

Protocol	Port	Description
TCP	1936	Metrics
	9000-9999	Host level services, including the node exporter on ports 9100-9101 and the Cluster Version Operator on port 9099 .
	10250-10259	The default ports that Kubernetes reserves
UDP	4789	VXLAN
	6081	Geneve
	9000-9999	Host level services, including the node exporter on ports 9100-9101 .
	500	IPsec IKE packets
	4500	IPsec NAT-T packets
	123	Network Time Protocol (NTP) on UDP port 123 If an external NTP time server is configured, you must open UDP port 123 .
TCP/UDP	30000-32767	Kubernetes node port
ESP	N/A	IPsec Encapsulating Security Payload (ESP)

Table 18.60. Ports used for all-machine to control plane communications

Protocol	Port	Description
TCP	6443	Kubernetes API

Table 18.61. Ports used for control plane machine to control plane machine communications

Protocol	Port	Description
TCP	2379-2380	etcd server and peer ports

NTP configuration for user-provisioned infrastructure

OpenShift Container Platform clusters are configured to use a public Network Time Protocol (NTP) server by default. If you want to use a local enterprise NTP server, or if your cluster is being deployed in a disconnected network, you can configure the cluster to use a specific time server. For more information, see the documentation for *Configuring chrony time service*.

If a DHCP server provides NTP server information, the chrony time service on the Red Hat Enterprise Linux CoreOS (RHCOS) machines read the information and can sync the clock with the NTP servers.

Additional resources

- [Configuring chrony time service](#)

18.5.4.11. User-provisioned DNS requirements

In OpenShift Container Platform deployments, DNS name resolution is required for the following components:


- The Kubernetes API
- The OpenShift Container Platform application wildcard
- The bootstrap, control plane, and compute machines

Reverse DNS resolution is also required for the Kubernetes API, the bootstrap machine, the control plane machines, and the compute machines.

DNS A/AAAA or CNAME records are used for name resolution and PTR records are used for reverse name resolution. The reverse records are important because Red Hat Enterprise Linux CoreOS (RHCOS) uses the reverse records to set the hostnames for all the nodes, unless the hostnames are provided by DHCP. Additionally, the reverse records are used to generate the certificate signing requests (CSR) that OpenShift Container Platform needs to operate.

The following DNS records are required for a user-provisioned OpenShift Container Platform cluster and they must be in place before installation. In each record, **<cluster_name>** is the cluster name and **<base_domain>** is the base domain that you specify in the **install-config.yaml** file. A complete DNS record takes the form: **<component>.<cluster_name>.<base_domain>..**

Table 18.62. Required DNS records

Component	Record	Description
Kubernetes API	api.<cluster_name>.<base_domain>..	A DNS A/AAAA or CNAME record, and a DNS PTR record, to identify the API load balancer. These records must be resolvable by both clients external to the cluster and from all the nodes within the cluster.
	api-int.<cluster_name>.<base_domain>..	A DNS A/AAAA or CNAME record, and a DNS PTR record, to internally identify the API load balancer. These records must be resolvable from all the nodes within the cluster.
		 <p>IMPORTANT</p> <p>The API server must be able to resolve the worker nodes by the hostnames that are recorded in Kubernetes. If the API server cannot resolve the node names, then proxied API calls can fail, and you cannot retrieve logs from pods.</p>

Component	Record	Description
Routes	*.apps.<cluster_name>.<base_domain>.	A wildcard DNS A/AAAA or CNAME record that refers to the application ingress load balancer. The application ingress load balancer targets the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default. These records must be resolvable by both clients external to the cluster and from all the nodes within the cluster. For example, console-openshift-console.apps.<cluster_name>.<base_domain> is used as a wildcard route to the OpenShift Container Platform console.
Bootstrap machine	bootstrap.<cluster_name>.<base_domain>.	A DNS A/AAAA or CNAME record, and a DNS PTR record, to identify the bootstrap machine. These records must be resolvable by the nodes within the cluster.
Control plane machines	<control_plane><n>.<cluster_name>.<base_domain>.	DNS A/AAAA or CNAME records and DNS PTR records to identify each machine for the control plane nodes. These records must be resolvable by the nodes within the cluster.
Compute machines	<compute><n>.<cluster_name>.<base_domain>.	DNS A/AAAA or CNAME records and DNS PTR records to identify each machine for the worker nodes. These records must be resolvable by the nodes within the cluster.



NOTE

In OpenShift Container Platform 4.4 and later, you do not need to specify etcd host and SRV records in your DNS configuration.

TIP

You can use the **dig** command to verify name and reverse name resolution. See the section on *Validating DNS resolution for user-provisioned infrastructure* for detailed validation steps.

18.5.4.11.1. Example DNS configuration for user-provisioned clusters

This section provides A and PTR record configuration samples that meet the DNS requirements for deploying OpenShift Container Platform on user-provisioned infrastructure. The samples are not meant to provide advice for choosing one DNS solution over another.

In the examples, the cluster name is **ocp4** and the base domain is **example.com**.

Example DNS A record configuration for a user-provisioned cluster

The following example is a BIND zone file that shows sample A records for name resolution in a user-provisioned cluster.


Example 18.10. Sample DNS zone database

```

$TTL 1W
@ IN SOA ns1.example.com. root (
    2019070700 ; serial
    3H ; refresh (3 hours)
    30M ; retry (30 minutes)
    2W ; expiry (2 weeks)
    1W ) ; minimum (1 week)
IN NS ns1.example.com.
IN MX 10 smtp.example.com.
;
;
;
ns1.example.com. IN A 192.168.1.5
smtp.example.com. IN A 192.168.1.5
;
;
helper.example.com. IN A 192.168.1.5
helper.ocp4.example.com. IN A 192.168.1.5
;
;
api.ocp4.example.com. IN A 192.168.1.5 1
api-int.ocp4.example.com. IN A 192.168.1.5 2
;
;
*.apps.ocp4.example.com. IN A 192.168.1.5 3
;
;
bootstrap.ocp4.example.com. IN A 192.168.1.96 4
;
;
control-plane0.ocp4.example.com. IN A 192.168.1.97 5
control-plane1.ocp4.example.com. IN A 192.168.1.98 6
control-plane2.ocp4.example.com. IN A 192.168.1.99 7
;
;
compute0.ocp4.example.com. IN A 192.168.1.11 8
compute1.ocp4.example.com. IN A 192.168.1.7 9
;
;
;EOF

```

- 1 Provides name resolution for the Kubernetes API. The record refers to the IP address of the API load balancer.
- 2 Provides name resolution for the Kubernetes API. The record refers to the IP address of the API load balancer and is used for internal cluster communications.
- 3 Provides name resolution for the wildcard routes. The record refers to the IP address of the application ingress load balancer. The application ingress load balancer targets the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.

**NOTE**

In the example, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

- 4 Provides name resolution for the bootstrap machine.
- 5 6 7 Provides name resolution for the control plane machines.
- 8 9 Provides name resolution for the compute machines.

Example DNS PTR record configuration for a user-provisioned cluster

The following example BIND zone file shows sample PTR records for reverse name resolution in a user-provisioned cluster.

Example 18.11. Sample DNS zone database for reverse records

```
$TTL 1W
@ IN SOA ns1.example.com. root (
    2019070700 ; serial
    3H ; refresh (3 hours)
    30M ; retry (30 minutes)
    2W ; expiry (2 weeks)
    1W ) ; minimum (1 week)
IN NS ns1.example.com.
;
5.1.168.192.in-addr.arpa. IN PTR api.ocp4.example.com. 1
5.1.168.192.in-addr.arpa. IN PTR api-int.ocp4.example.com. 2
;
96.1.168.192.in-addr.arpa. IN PTR bootstrap.ocp4.example.com. 3
;
97.1.168.192.in-addr.arpa. IN PTR control-plane0.ocp4.example.com. 4
98.1.168.192.in-addr.arpa. IN PTR control-plane1.ocp4.example.com. 5
99.1.168.192.in-addr.arpa. IN PTR control-plane2.ocp4.example.com. 6
;
11.1.168.192.in-addr.arpa. IN PTR compute0.ocp4.example.com. 7
7.1.168.192.in-addr.arpa. IN PTR compute1.ocp4.example.com. 8
;
;EOF
```

- 1 Provides reverse DNS resolution for the Kubernetes API. The PTR record refers to the record name of the API load balancer.
- 2 Provides reverse DNS resolution for the Kubernetes API. The PTR record refers to the record name of the API load balancer and is used for internal cluster communications.
- 3 Provides reverse DNS resolution for the bootstrap machine.
- 4 5 6 Provides reverse DNS resolution for the control plane machines.
- 7 8 Provides reverse DNS resolution for the compute machines.

**NOTE**

A PTR record is not required for the OpenShift Container Platform application wildcard.

18.5.4.12. Load balancing requirements for user-provisioned infrastructure

Before you install OpenShift Container Platform, you must provision the API and application Ingress load balancing infrastructure. In production scenarios, you can deploy the API and application Ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

**NOTE**

If you want to deploy the API and application Ingress load balancers with a Red Hat Enterprise Linux (RHEL) instance, you must purchase the RHEL subscription separately.

The load balancing infrastructure must meet the following requirements:

1. **API load balancer.** Provides a common endpoint for users, both human and machine, to interact with and configure the platform. Configure the following conditions:
 - Layer 4 load balancing only. This can be referred to as Raw TCP or SSL Passthrough mode.
 - A stateless load balancing algorithm. The options vary based on the load balancer implementation.

**IMPORTANT**

Do not configure session persistence for an API load balancer. Configuring session persistence for a Kubernetes API server might cause performance issues from excess application traffic for your OpenShift Container Platform cluster and the Kubernetes API that runs inside the cluster.

Configure the following ports on both the front and back of the load balancers:

Table 18.63. API load balancer

Port	Back-end machines (pool members)	Internal	External	Description
6443	Bootstrap and control plane. You remove the bootstrap machine from the load balancer after the bootstrap machine initializes the cluster control plane. You must configure the /readyz endpoint for the API server health check probe.	X	X	Kubernetes API server
22623	Bootstrap and control plane. You remove the bootstrap machine from the load balancer after the bootstrap machine initializes the cluster control plane.	X		Machine config server

**NOTE**

The load balancer must be configured to take a maximum of 30 seconds from the time the API server turns off the **/readyz** endpoint to the removal of the API server instance from the pool. Within the time frame after **/readyz** returns an error or becomes healthy, the endpoint must have been removed or added. Probing every 5 or 10 seconds, with two successful requests to become healthy and three to become unhealthy, are well-tested values.

2. **Application Ingress load balancer.** Provides an ingress point for application traffic flowing in from outside the cluster. A working configuration for the Ingress router is required for an OpenShift Container Platform cluster.

Configure the following conditions:

- Layer 4 load balancing only. This can be referred to as Raw TCP or SSL Passthrough mode.
- A connection-based or session-based persistence is recommended, based on the options available and types of applications that will be hosted on the platform.

TIP

If the true IP address of the client can be seen by the application Ingress load balancer, enabling source IP-based session persistence can improve performance for applications that use end-to-end TLS encryption.

Configure the following ports on both the front and back of the load balancers:

Table 18.64. Application Ingress load balancer

Port	Back-end machines (pool members)	Internal	External	Description
443	The machines that run the Ingress Controller pods, compute, or worker, by default.	X	X	HTTPS traffic
80	The machines that run the Ingress Controller pods, compute, or worker, by default.	X	X	HTTP traffic

**NOTE**

If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application Ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes.

18.5.4.12.1. Example load balancer configuration for user-provisioned clusters

This section provides an example API and application Ingress load balancer configuration that meets the load balancing requirements for user-provisioned clusters. The sample is an **/etc/haproxy/haproxy.cfg** configuration for an HAProxy load balancer. The example is not meant to provide advice for choosing one load balancing solution over another.

In the example, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.



NOTE

If you are using HAProxy as a load balancer and SELinux is set to **enforcing**, you must ensure that the HAProxy service can bind to the configured TCP port by running **setsebool -P haproxy_connect_any=1**.

Example 18.12. Sample API and application Ingress load balancer configuration

```

global
  log      127.0.0.1 local2
  pidfile  /var/run/haproxy.pid
  maxconn  4000
  daemon
defaults
  mode            http
  log             global
  option          dontlognull
  option http-server-close
  option          redispatch
  retries         3
  timeout http-request 10s
  timeout queue    1m
  timeout connect  10s
  timeout client   1m
  timeout server   1m
  timeout http-keep-alive 10s
  timeout check    10s
  maxconn          3000
listen api-server-6443 1
  bind *:6443
  mode tcp
  option httpchk GET /readyz HTTP/1.0
  option log-health-checks
  balance roundrobin
  server bootstrap bootstrap.ocp4.example.com:6443 verify none check check-ssl inter 10s fall 2
  rise 3 backup 2
  server master0 master0.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
  fall 2 rise 3
  server master1 master1.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
  fall 2 rise 3
  server master2 master2.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
  fall 2 rise 3
listen machine-config-server-22623 3
  bind *:22623
  mode tcp
  server bootstrap bootstrap.ocp4.example.com:22623 check inter 1s backup 4
  server master0 master0.ocp4.example.com:22623 check inter 1s
  server master1 master1.ocp4.example.com:22623 check inter 1s
  server master2 master2.ocp4.example.com:22623 check inter 1s
listen ingress-router-443 5

```




```

bind *:443
mode tcp
balance source
server compute0 compute0.ocp4.example.com:443 check inter 1s
server compute1 compute1.ocp4.example.com:443 check inter 1s
listen ingress-router-80 6
bind *:80
mode tcp
balance source
server compute0 compute0.ocp4.example.com:80 check inter 1s
server compute1 compute1.ocp4.example.com:80 check inter 1s

```

- 1** Port **6443** handles the Kubernetes API traffic and points to the control plane machines.
- 2** **4** The bootstrap entries must be in place before the OpenShift Container Platform cluster installation and they must be removed after the bootstrap process is complete.
- 3** Port **22623** handles the machine config server traffic and points to the control plane machines.
- 5** Port **443** handles the HTTPS traffic and points to the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.
- 6** Port **80** handles the HTTP traffic and points to the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.



NOTE

If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application Ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes.

TIP

If you are using HAProxy as a load balancer, you can check that the **haproxy** process is listening on ports **6443**, **22623**, **443**, and **80** by running **netstat -nltp** on the HAProxy node.

18.5.5. Preparing the user-provisioned infrastructure

Before you install OpenShift Container Platform on user-provisioned infrastructure, you must prepare the underlying infrastructure.

This section provides details about the high-level steps required to set up your cluster infrastructure in preparation for an OpenShift Container Platform installation. This includes configuring IP networking and network connectivity for your cluster nodes, enabling the required ports through your firewall, and setting up the required DNS and load balancing infrastructure.

After preparation, your cluster infrastructure must meet the requirements outlined in the *Requirements for a cluster with user-provisioned infrastructure* section.

Prerequisites

- You have reviewed the [OpenShift Container Platform 4.x Tested Integrations](#) page.
- You have reviewed the infrastructure requirements detailed in the *Requirements for a cluster with user-provisioned infrastructure* section.

Procedure

1. If you are using DHCP to provide the IP networking configuration to your cluster nodes, configure your DHCP service.
 - a. Add persistent IP addresses for the nodes to your DHCP server configuration. In your configuration, match the MAC address of the relevant network interface to the intended IP address for each node.
 - b. When you use DHCP to configure IP addressing for the cluster machines, the machines also obtain the DNS server information through DHCP. Define the persistent DNS server address that is used by the cluster nodes through your DHCP server configuration.



NOTE

If you are not using a DHCP service, you must provide the IP networking configuration and the address of the DNS server to the nodes at RHCOS install time. These can be passed as boot arguments if you are installing from an ISO image. See the *Installing RHCOS and starting the OpenShift Container Platform bootstrap process* section for more information about static IP provisioning and advanced networking options.

- c. Define the hostnames of your cluster nodes in your DHCP server configuration. See the *Setting the cluster node hostnames through DHCP* section for details about hostname considerations.



NOTE

If you are not using a DHCP service, the cluster nodes obtain their hostname through a reverse DNS lookup.

2. Choose to perform either a fast track installation of Red Hat Enterprise Linux CoreOS (RHCOS) or a full installation of Red Hat Enterprise Linux CoreOS (RHCOS). For the full installation, you must set up an HTTP or HTTPS server to provide Ignition files and install images to the cluster nodes. For the fast track installation an HTTP or HTTPS server is not required, however, a DHCP server is required. See sections "Fast-track installation: Creating Red Hat Enterprise Linux CoreOS (RHCOS) machines" and "Full installation: Creating Red Hat Enterprise Linux CoreOS (RHCOS) machines".
3. Ensure that your network infrastructure provides the required network connectivity between the cluster components. See the *Networking requirements for user-provisioned infrastructure* section for details about the requirements.
4. Configure your firewall to enable the ports required for the OpenShift Container Platform cluster components to communicate. See *Networking requirements for user-provisioned infrastructure* section for details about the ports that are required.



IMPORTANT

By default, port **1936** is accessible for an OpenShift Container Platform cluster, because each control plane node needs access to this port.

Avoid using the Ingress load balancer to expose this port, because doing so might result in the exposure of sensitive information, such as statistics and metrics, related to Ingress Controllers.

5. Setup the required DNS infrastructure for your cluster.
 - a. Configure DNS name resolution for the Kubernetes API, the application wildcard, the bootstrap machine, the control plane machines, and the compute machines.
 - b. Configure reverse DNS resolution for the Kubernetes API, the bootstrap machine, the control plane machines, and the compute machines.
See the *User-provisioned DNS requirements* section for more information about the OpenShift Container Platform DNS requirements.
6. Validate your DNS configuration.
 - a. From your installation node, run DNS lookups against the record names of the Kubernetes API, the wildcard routes, and the cluster nodes. Validate that the IP addresses in the responses correspond to the correct components.
 - b. From your installation node, run reverse DNS lookups against the IP addresses of the load balancer and the cluster nodes. Validate that the record names in the responses correspond to the correct components.
See the *Validating DNS resolution for user-provisioned infrastructure* section for detailed DNS validation steps.
7. Provision the required API and application ingress load balancing infrastructure. See the *Load balancing requirements for user-provisioned infrastructure* section for more information about the requirements.



NOTE

Some load balancing solutions require the DNS name resolution for the cluster nodes to be in place before the load balancing is initialized.

18.5.6. Validating DNS resolution for user-provisioned infrastructure

You can validate your DNS configuration before installing OpenShift Container Platform on user-provisioned infrastructure.



IMPORTANT

The validation steps detailed in this section must succeed before you install your cluster.

Prerequisites

- You have configured the required DNS records for your user-provisioned infrastructure.

Procedure

1. From your installation node, run DNS lookups against the record names of the Kubernetes API, the wildcard routes, and the cluster nodes. Validate that the IP addresses contained in the responses correspond to the correct components.
 - a. Perform a lookup against the Kubernetes API record name. Check that the result points to the IP address of the API load balancer:

```
$ dig +noall +answer @<nameserver_ip> api.<cluster_name>.<base_domain> 1
```

- 1 Replace **<nameserver_ip>** with the IP address of the nameserver, **<cluster_name>** with your cluster name, and **<base_domain>** with your base domain name.

Example output

```
api.ocp4.example.com. 604800 IN A 192.168.1.5
```

- b. Perform a lookup against the Kubernetes internal API record name. Check that the result points to the IP address of the API load balancer:

```
$ dig +noall +answer @<nameserver_ip> api-int.<cluster_name>.<base_domain>
```

Example output

```
api-int.ocp4.example.com. 604800 IN A 192.168.1.5
```

- c. Test an example ***.apps.<cluster_name>.<base_domain>** DNS wildcard lookup. All of the application wildcard lookups must resolve to the IP address of the application ingress load balancer:

```
$ dig +noall +answer @<nameserver_ip> random.apps.<cluster_name>.<base_domain>
```

Example output

```
random.apps.ocp4.example.com. 604800 IN A 192.168.1.5
```



NOTE

In the example outputs, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

You can replace **random** with another wildcard value. For example, you can query the route to the OpenShift Container Platform console:

```
$ dig +noall +answer @<nameserver_ip> console-openshift-console.apps.<cluster_name>.<base_domain>
```

Example output

```
console-openshift-console.apps.ocp4.example.com. 604800 IN A 192.168.1.5
```

-
- d. Run a lookup against the bootstrap DNS record name. Check that the result points to the IP address of the bootstrap node:

```
$ dig +noall +answer @<nameserver_ip> bootstrap.<cluster_name>.<base_domain>
```

Example output

```
bootstrap.ocp4.example.com. 604800 IN A 192.168.1.96
```

- e. Use this method to perform lookups against the DNS record names for the control plane and compute nodes. Check that the results correspond to the IP addresses of each node.
2. From your installation node, run reverse DNS lookups against the IP addresses of the load balancer and the cluster nodes. Validate that the record names contained in the responses correspond to the correct components.

- a. Perform a reverse lookup against the IP address of the API load balancer. Check that the response includes the record names for the Kubernetes API and the Kubernetes internal API:

```
$ dig +noall +answer @<nameserver_ip> -x 192.168.1.5
```

Example output

```
5.1.168.192.in-addr.arpa. 604800 IN PTR api-int.ocp4.example.com. 1
5.1.168.192.in-addr.arpa. 604800 IN PTR api.ocp4.example.com. 2
```

1 Provides the record name for the Kubernetes internal API.

2 Provides the record name for the Kubernetes API.



NOTE

A PTR record is not required for the OpenShift Container Platform application wildcard. No validation step is needed for reverse DNS resolution against the IP address of the application ingress load balancer.

- b. Perform a reverse lookup against the IP address of the bootstrap node. Check that the result points to the DNS record name of the bootstrap node:

```
$ dig +noall +answer @<nameserver_ip> -x 192.168.1.96
```

Example output

```
96.1.168.192.in-addr.arpa. 604800 IN PTR bootstrap.ocp4.example.com.
```

- c. Use this method to perform reverse lookups against the IP addresses for the control plane and compute nodes. Check that the results correspond to the DNS record names of each node.

18.5.7. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.



NOTE

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.


2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.

**NOTE**

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```

**NOTE**

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

18.5.8. Manually creating the installation configuration file

Installing the cluster requires that you manually create the installation configuration file.


Prerequisites

- You have an SSH public key on your local machine to provide to the installation program. The key will be used for SSH authentication onto your cluster nodes for debugging and disaster recovery.
- You have obtained the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create an installation directory to store your required installation assets in:

```
$ mkdir <installation_directory>
```



IMPORTANT

You must create a directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.


2. Customize the sample **install-config.yaml** file template that is provided and save it in the **<installation_directory>**.



NOTE

You must name this configuration file **install-config.yaml**.

3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the next step of the installation process. You must back it up now.

Additional resources

- [Installation configuration parameters for IBM Z®](#)

18.5.8.1. Sample install-config.yaml file for IBM Z

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.

```
apiVersion: v1
baseDomain: example.com 1
compute: 2
- hyperthreading: Enabled 3
  name: worker
  replicas: 0 4
  architecture: s390x
controlPlane: 5
  hyperthreading: Enabled 6
  name: master
  replicas: 3 7
  architecture: s390x
metadata:
  name: test 8
networking:
  clusterNetwork:
  - cidr: 10.128.0.0/14 9
```



```


hostPrefix: 23 10
networkType: OVNKubernetes 11
serviceNetwork: 12
- 172.30.0.0/16
platform:
  none: {} 13
fips: false 14
pullSecret: '{"auths":{"<local_registry>":{"auth": "<credentials>","email": "you@example.com"}}}' 15
sshKey: 'ssh-ed25519 AAAA...' 16
additionalTrustBundle: | 17
-----BEGIN CERTIFICATE-----
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
-----END CERTIFICATE-----
imageContentSources: 18
- mirrors:
- <local_repository>/ocp4/openshift4
  source: quay.io/openshift-release-dev/ocp-release
- mirrors:
- <local_repository>/ocp4/openshift4
  source: quay.io/openshift-release-dev/ocp-v4.0-art-dev

```

1 The base domain of the cluster. All DNS records must be sub-domains of this base and include the cluster name.

2 **5** The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.

3 **6** Specifies whether to enable or disable simultaneous multithreading (SMT), or hyperthreading. By default, SMT is enabled to increase the performance of the cores in your machines. You can disable it by setting the parameter value to **Disabled**. If you disable SMT, you must disable it in all cluster machines; this includes both control plane and compute machines.



NOTE

Simultaneous multithreading (SMT) is enabled by default. If SMT is not available on your OpenShift Container Platform nodes, the **hyperthreading** parameter has no effect.



IMPORTANT


If you disable **hyperthreading**, whether on your OpenShift Container Platform nodes or in the **install-config.yaml** file, ensure that your capacity planning accounts for the dramatically decreased machine performance.

4 You must set this value to **0** when you install OpenShift Container Platform on user-provisioned infrastructure. In installer-provisioned installations, the parameter controls the number of compute machines that the cluster creates and manages for you. In user-provisioned installations, you must manually deploy the compute machines before you finish installing the cluster.

**NOTE**

If you are installing a three-node cluster, do not deploy any compute machines when you install the Red Hat Enterprise Linux CoreOS (RHCOS) machines.

- 7 The number of control plane machines that you add to the cluster. Because the cluster uses these values as the number of etcd endpoints in the cluster, the value must match the number of control plane machines that you deploy.
- 8 The cluster name that you specified in your DNS records.
- 9 A block of IP addresses from which pod IP addresses are allocated. This block must not overlap with existing physical networks. These IP addresses are used for the pod network. If you need to access the pods from an external network, you must configure load balancers and routers to manage the traffic.

**NOTE**

Class E CIDR range is reserved for a future use. To use the Class E CIDR range, you must ensure your networking environment accepts the IP addresses within the Class E CIDR range.

- 10 The subnet prefix length to assign to each individual node. For example, if **hostPrefix** is set to **23**, then each node is assigned a **/23** subnet out of the given **cidr**, which allows for 510 ($2^{(32 - 23)} - 2$) pod IP addresses. If you are required to provide access to nodes from an external network, configure load balancers and routers to manage the traffic.
- 11 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 12 The IP address pool to use for service IP addresses. You can enter only one IP address pool. This block must not overlap with existing physical networks. If you need to access the services from an external network, configure load balancers and routers to manage the traffic.
- 13 You must set the platform to **none**. You cannot provide additional platform configuration variables for IBM Z[®] infrastructure.

**IMPORTANT**

Clusters that are installed with the platform type **none** are unable to use some features, such as managing compute machines with the Machine API. This limitation applies even if the compute machines that are attached to the cluster are installed on a platform that would normally support the feature. This parameter cannot be changed after installation.

- 14 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.



IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 15 For **<local_registry>**, specify the registry domain name, and optionally the port, that your mirror registry uses to serve content. For example, **registry.example.com** or **registry.example.com:5000**. For **<credentials>**, specify the base64-encoded user name and password for your mirror registry.
- 16 The SSH public key for the **core** user in Red Hat Enterprise Linux CoreOS (RHCOS).



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- 17 Add the **additionalTrustBundle** parameter and value. The value must be the contents of the certificate file that you used for your mirror registry. The certificate file can be an existing, trusted certificate authority or the self-signed certificate that you generated for the mirror registry.
- 18 Provide the **imageContentSources** section according to the output of the command that you used to mirror the repository.



IMPORTANT

- When using the **oc adm release mirror** command, use the output from the **imageContentSources** section.
- When using **oc mirror** command, use the **repositoryDigestMirrors** section of the **ImageContentSourcePolicy** file that results from running the command.
- **ImageContentSourcePolicy** is deprecated. For more information see *Configuring image registry repository mirroring*.

18.5.8.2. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).


Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when

http/https proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.

**NOTE**

The installation program does not support the proxy **readinessEndpoints** field.

**NOTE**

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.

**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

18.5.8.3. Configuring a three-node cluster

Optionally, you can deploy zero compute machines in a minimal three node cluster that consists of three control plane machines only. This provides smaller, more resource efficient clusters for cluster administrators and developers to use for testing, development, and production.

In three-node OpenShift Container Platform environments, the three control plane machines are schedulable, which means that your application workloads are scheduled to run on them.


Prerequisites

- You have an existing **install-config.yaml** file.

Procedure

- Ensure that the number of compute replicas is set to **0** in your **install-config.yaml** file, as shown in the following **compute** stanza:

```
compute:
- name: worker
  platform: {}
  replicas: 0
```

**NOTE**

You must set the value of the **replicas** parameter for the compute machines to **0** when you install OpenShift Container Platform on user-provisioned infrastructure, regardless of the number of compute machines you are deploying. In installer-provisioned installations, the parameter controls the number of compute machines that the cluster creates and manages for you. This does not apply to user-provisioned installations, where the compute machines are deployed manually.

**NOTE**

The preferred resource for control plane nodes is six vCPUs and 21 GB. For three control plane nodes this is the memory + vCPU equivalent of a minimum five-node cluster. You should back the three nodes, each installed on a 120 GB disk, with three IFLs that are SMT2 enabled. The minimum tested setup is three vCPUs and 10 GB on a 120 GB disk for each control plane node.

For three-node cluster installations, follow these next steps:

- If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes. See the *Load balancing requirements for user-provisioned infrastructure* section for more information.
- When you create the Kubernetes manifest files in the following procedure, ensure that the **mastersSchedulable** parameter in the `<installation_directory>/manifests/cluster-scheduler-02-config.yml` file is set to **true**. This enables your application workloads to run on the control plane nodes.
- Do not deploy any compute nodes when you create the Red Hat Enterprise Linux CoreOS (RHCOS) machines.

18.5.9. Cluster Network Operator configuration

The configuration for the cluster network is specified as part of the Cluster Network Operator (CNO) configuration and stored in a custom resource (CR) object that is named **cluster**. The CR specifies the fields for the **Network** API in the **operator.openshift.io** API group.

The CNO configuration inherits the following fields during cluster installation from the **Network** API in the **Network.config.openshift.io** API group:

clusterNetwork

IP address pools from which pod IP addresses are allocated.

serviceNetwork

IP address pool for services.

defaultNetwork.type

Cluster network plugin. **OVNKubernetes** is the only supported plugin during installation.

You can specify the cluster network plugin configuration for your cluster by setting the fields for the **defaultNetwork** object in the CNO object named **cluster**.

18.5.9.1. Cluster Network Operator configuration object

The fields for the Cluster Network Operator (CNO) are described in the following table:

Table 18.65. Cluster Network Operator configuration object


Field	Type	Description
metadata.name	string	The name of the CNO object. This name is always cluster .
spec.clusterNetwork	array	<p>A list specifying the blocks of IP addresses from which pod IP addresses are allocated and the subnet prefix length assigned to each individual node in the cluster. For example:</p> <pre>spec: clusterNetwork: - cidr: 10.128.0.0/19 hostPrefix: 23 - cidr: 10.128.32.0/19 hostPrefix: 23</pre>
spec.serviceNetwork	array	<p>A block of IP addresses for services. The OpenShift SDN and OVN-Kubernetes network plugins support only a single IP address block for the service network. For example:</p> <pre>spec: serviceNetwork: - 172.30.0.0/14</pre> <p>You can customize this field only in the install-config.yaml file before you create the manifests. The value is read-only in the manifest file.</p>
spec.defaultNetwork	object	Configures the network plugin for the cluster network.
spec.kubeProxy Config	object	The fields for this object specify the kube-proxy configuration. If you are using the OVN-Kubernetes cluster network plugin, the kube-proxy configuration has no effect.

defaultNetwork object configuration

The values for the **defaultNetwork** object are defined in the following table:

Table 18.66. defaultNetwork object

Field	Type	Description
-------	------	-------------

Field	Type	Description
type	string	<p>OVNKubernetes. The Red Hat OpenShift Networking network plugin is selected during installation. This value cannot be changed after cluster installation.</p>  <p>NOTE</p> <p>OpenShift Container Platform uses the OVN-Kubernetes network plugin by default. OpenShift SDN is no longer available as an installation choice for new clusters.</p>
ovnKubernetesConfig	object	This object is only valid for the OVN-Kubernetes network plugin.

Configuration for the OVN-Kubernetes network plugin

The following table describes the configuration fields for the OVN-Kubernetes network plugin:

Table 18.67. ovnKubernetesConfig object

Field	Type	Description
mtu	integer	<p>The maximum transmission unit (MTU) for the Geneve (Generic Network Virtualization Encapsulation) overlay network. This is detected automatically based on the MTU of the primary network interface. You do not normally need to override the detected MTU.</p> <p>If the auto-detected value is not what you expect it to be, confirm that the MTU on the primary network interface on your nodes is correct. You cannot use this option to change the MTU value of the primary network interface on the nodes.</p> <p>If your cluster requires different MTU values for different nodes, you must set this value to 100 less than the lowest MTU value in your cluster. For example, if some nodes in your cluster have an MTU of 9001, and some have an MTU of 1500, you must set this value to 1400.</p>
genevePort	integer	The port to use for all Geneve packets. The default value is 6081 . This value cannot be changed after cluster installation.
ipsecConfig	object	Specify a configuration object for customizing the IPsec configuration.
ipv4	object	Specifies a configuration object for IPv4 settings.


Field	Type	Description
ipv6	object	Specifies a configuration object for IPv6 settings.
policyAuditConfig	object	Specify a configuration object for customizing network policy audit logging. If unset, the defaults audit log settings are used.
gatewayConfig	object	Optional: Specify a configuration object for customizing how egress traffic is sent to the node gateway.  <p>NOTE</p> <p>While migrating egress traffic, you can expect some disruption to workloads and service traffic until the Cluster Network Operator (CNO) successfully rolls out the changes.</p>

Table 18.68. `ovnKubernetesConfig.ipv4` object

Field	Type	Description
internalTransitSwitchSubnet	string	If your existing network infrastructure overlaps with the 100.88.0.0/16 IPv4 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. The subnet for the distributed transit switch that enables east-west traffic. This subnet cannot overlap with any other subnets used by OVN-Kubernetes or on the host itself. It must be large enough to accommodate one IP address per node in your cluster. The default value is 100.88.0.0/16 .
internalJoinSubnet	string	If your existing network infrastructure overlaps with the 100.64.0.0/16 IPv4 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. You must ensure that the IP address range does not overlap with any other subnet used by your OpenShift Container Platform installation. The IP address range must be larger than the maximum number of nodes that can be added to the cluster. For example, if the clusterNetwork.cidr value is 10.128.0.0/14 and the clusterNetwork.hostPrefix value is /23 , then the maximum number of nodes is 2⁽²³⁻¹⁴⁾=512 . The default value is 100.64.0.0/16 .

Table 18.69. `ovnKubernetesConfig.ipv6` object

Field	Type	Description
-------	------	-------------

Field	Type	Description
internalTransitSwitchSubnet	string	<p>If your existing network infrastructure overlaps with the fd97::/64 IPv6 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. The subnet for the distributed transit switch that enables east-west traffic. This subnet cannot overlap with any other subnets used by OVN-Kubernetes or on the host itself. It must be large enough to accommodate one IP address per node in your cluster.</p> <p>The default value is fd97::/64.</p>
internalJoinSubnet	string	<p>If your existing network infrastructure overlaps with the fd98::/64 IPv6 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. You must ensure that the IP address range does not overlap with any other subnet used by your OpenShift Container Platform installation. The IP address range must be larger than the maximum number of nodes that can be added to the cluster.</p> <p>The default value is fd98::/64.</p>

Table 18.70. `policyAuditConfig` object

Field	Type	Description
rateLimit	integer	The maximum number of messages to generate every second per node. The default value is 20 messages per second.
maxFileSize	integer	The maximum size for the audit log in bytes. The default value is 50000000 or 50 MB.
maxLogFiles	integer	The maximum number of log files that are retained.
destination	string	<p>One of the following additional audit log targets:</p> <p>libc The libc syslog() function of the journald process on the host.</p> <p>udp:<host>:<port> A syslog server. Replace <host>:<port> with the host and port of the syslog server.</p> <p>unix:<file> A Unix Domain Socket file specified by <file>.</p> <p>null Do not send the audit logs to any additional target.</p>

Field	Type	Description
syslogFacility	string	The syslog facility, such as kern , as defined by RFC5424. The default value is local0 .

Table 18.71. gatewayConfig object

Field	Type	Description
routingViaHost	boolean	Set this field to true to send egress traffic from pods to the host networking stack. For highly-specialized installations and applications that rely on manually configured routes in the kernel routing table, you might want to route egress traffic to the host networking stack. By default, egress traffic is processed in OVN to exit the cluster and is not affected by specialized routes in the kernel routing table. The default value is false . This field has an interaction with the Open vSwitch hardware offloading feature. If you set this field to true , you do not receive the performance benefits of the offloading because egress traffic is processed by the host networking stack.
ipForwarding	object	You can control IP forwarding for all traffic on OVN-Kubernetes managed interfaces by using the ipForwarding specification in the Network resource. Specify Restricted to only allow IP forwarding for Kubernetes related traffic. Specify Global to allow forwarding of all IP traffic. For new installations, the default is Restricted . For updates to OpenShift Container Platform 4.14 or later, the default is Global .
ipv4	object	Optional: Specify an object to configure the internal OVN-Kubernetes masquerade address for host to service traffic for IPv4 addresses.
ipv6	object	Optional: Specify an object to configure the internal OVN-Kubernetes masquerade address for host to service traffic for IPv6 addresses.

Table 18.72. gatewayConfig.ipv4 object

Field	Type	Description
internalMasqueradeSubnet	string	The masquerade IPv4 addresses that are used internally to enable host to service traffic. The host is configured with these IP addresses as well as the shared gateway bridge interface. The default value is 169.254.169.0/29 .

Table 18.73. gatewayConfig.ipv6 object


Field	Type	Description
internalMasqueradeSubnet	string	The masquerade IPv6 addresses that are used internally to enable host to service traffic. The host is configured with these IP addresses as well as the shared gateway bridge interface. The default value is fd69::/125 .

Table 18.74. ipsecConfig object

Field	Type	Description
mode	string	Specifies the behavior of the IPsec implementation. Must be one of the following values: <ul style="list-style-type: none"> ● Disabled: IPsec is not enabled on cluster nodes. ● External: IPsec is enabled for network traffic with external hosts. ● Full: IPsec is enabled for pod traffic and network traffic with external hosts.

Example OVN-Kubernetes configuration with IPsec enabled

```
defaultNetwork:
  type: OVNKubernetes
  ovnKubernetesConfig:
    mtu: 1400
    genevePort: 6081
    ipsecConfig:
      mode: Full
```



IMPORTANT

Using OVNKubernetes can lead to a stack exhaustion problem on IBM Power®.

kubeProxyConfig object configuration (OpenShiftSDN container network interface only)

The values for the **kubeProxyConfig** object are defined in the following table:

Table 18.75. kubeProxyConfig object

Field	Type	Description
-------	------	-------------

Field	Type	Description
iptablesSyncPeriod	string	<p>The refresh period for iptables rules. The default value is 30s. Valid suffixes include s, m, and h and are described in the Go time package documentation.</p> <div style="display: flex; align-items: flex-start;"> <div style="width: 40px; height: 100px; background: repeating-linear-gradient(45deg, transparent, transparent 2px, #ccc 2px, #ccc 4px); margin-right: 10px;"></div> <div> <p>NOTE</p> <p>Because of performance improvements introduced in OpenShift Container Platform 4.3 and greater, adjusting the iptablesSyncPeriod parameter is no longer necessary.</p> </div> </div>
proxyArguments.iptables-min-sync-period	array	<p>The minimum duration before refreshing iptables rules. This field ensures that the refresh does not happen too frequently. Valid suffixes include s, m, and h and are described in the Go time package. The default value is:</p> <pre>kubeProxyConfig: proxyArguments: iptables-min-sync-period: - 0s</pre>

18.5.10. Creating the Kubernetes manifest and Ignition config files

Because you must modify some cluster definition files and manually start the cluster machines, you must generate the Kubernetes manifest and Ignition config files that the cluster needs to configure the machines.

The installation configuration file transforms into the Kubernetes manifests. The manifests wrap into the Ignition configuration files, which are later used to configure the cluster machines.



IMPORTANT

- The Ignition config files that the OpenShift Container Platform installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.



NOTE

The installation program that generates the manifest and Ignition files is architecture specific and can be obtained from the [client image mirror](#). The Linux version of the installation program runs on s390x only. This installer program is also available as a Mac OS version.

Prerequisites

- You obtained the OpenShift Container Platform installation program. For a restricted network installation, these files are on your mirror host.
- You created the **install-config.yaml** installation configuration file.

Procedure

1. Change to the directory that contains the OpenShift Container Platform installation program and generate the Kubernetes manifests for the cluster:

```
$ ./openshift-install create manifests --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the installation directory that contains the **install-config.yaml** file you created.



WARNING

If you are installing a three-node cluster, skip the following step to allow the control plane nodes to be schedulable.



IMPORTANT

When you configure control plane nodes from the default unschedulable to schedulable, additional subscriptions are required. This is because control plane nodes then become compute nodes.

2. Check that the **mastersSchedulable** parameter in the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** Kubernetes manifest file is set to **false**. This setting prevents pods from being scheduled on the control plane machines:
 - a. Open the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** file.
 - b. Locate the **mastersSchedulable** parameter and ensure that it is set to **false**.
 - c. Save and exit the file.
3. To create the Ignition configuration files, run the following command from the directory that contains the installation program:

```
$ ./openshift-install create ignition-configs --dir <installation_directory> 1
```

1 For **<installation_directory>**, specify the same installation directory.

Ignition config files are created for the bootstrap, control plane, and compute nodes in the installation directory. The **kubeadmin-password** and **kubeconfig** files are created in the **./<installation_directory>/auth** directory:

```
.
├── auth
│   ├── kubeadmin-password
│   └── kubeconfig
├── bootstrap.ign
├── master.ign
├── metadata.json
└── worker.ign
```

18.5.11. Installing RHCOS and starting the OpenShift Container Platform bootstrap process

To install OpenShift Container Platform on IBM Z® infrastructure that you provision, you must install Red Hat Enterprise Linux CoreOS (RHCOS) as Red Hat Enterprise Linux (RHEL) guest virtual machines. When you install RHCOS, you must provide the Ignition config file that was generated by the OpenShift Container Platform installation program for the type of machine you are installing. If you have configured suitable networking, DNS, and load balancing infrastructure, the OpenShift Container Platform bootstrap process begins automatically after the RHCOS machines have rebooted.

You can perform a fast-track installation of RHCOS that uses a prepackaged QEMU copy-on-write (QCOW2) disk image. Alternatively, you can perform a full installation on a new QCOW2 disk image.

To add further security to your system, you can optionally install RHCOS using IBM® Secure Execution before proceeding to the fast-track installation.

18.5.11.1. Installing RHCOS using IBM Secure Execution

Before you install RHCOS using IBM® Secure Execution, you must prepare the underlying infrastructure.

Prerequisites

- IBM® z15 or later, or IBM® LinuxONE III or later.
- Red Hat Enterprise Linux (RHEL) 8 or later.
- You have a bootstrap Ignition file. The file is not protected, enabling others to view and edit it.
- You have verified that the boot image has not been altered after installation.
- You must run all your nodes as IBM® Secure Execution guests.

Procedure

1. Prepare your RHEL KVM host to support IBM® Secure Execution.
 - By default, KVM hosts do not support guests in IBM® Secure Execution mode. To support guests in IBM® Secure Execution mode, KVM hosts must boot in LPAR mode with the kernel parameter specification **prot_virt=1**. To enable **prot_virt=1** on RHEL 8, follow these steps:
 - a. Navigate to **/boot/loader/entries/** to modify your bootloader configuration file ***.conf**.
 - b. Add the kernel command line parameter **prot_virt=1**.
 - c. Run the **zipl** command and reboot your system.

KVM hosts that successfully start with support for IBM® Secure Execution for Linux issue the following kernel message:

```
prot_virt: Reserving <amount>MB as ultravisor base storage.
```

- d. To verify that the KVM host now supports IBM® Secure Execution, run the following command:

```
# cat /sys/firmware/uv/prot_virt_host
```

Example output

```
1
```

The value of this attribute is 1 for Linux instances that detect their environment as consistent with that of a secure host. For other instances, the value is 0.

2. Add your host keys to the KVM guest via Ignition.

During the first boot, RHCOS looks for your host keys to re-encrypt itself with them. RHCOS searches for files starting with **ibm-z-hostkey-** in the **/etc/se-hostkeys** directory. All host keys, for each machine the cluster is running on, must be loaded into the directory by the administrator. After first boot, you cannot run the VM on any other machines.

**NOTE**

You need to prepare your Ignition file on a safe system. For example, another IBM® Secure Execution guest.

For example:

```
{
  "ignition": { "version": "3.0.0" },
  "storage": {
    "files": [
      {
        "path": "/etc/se-hostkeys/ibm-z-hostkey-<your-hostkey>.crt",
        "contents": {
          "source": "data:;base64,<base64 encoded hostkey document>"
        },
        "mode": 420
      },
      {
        "path": "/etc/se-hostkeys/ibm-z-hostkey-<your-hostkey>.crt",
        "contents": {
          "source": "data:;base64,<base64 encoded hostkey document>"
        },
        "mode": 420
      }
    ]
  }
}
```

**NOTE**

You can add as many host keys as required if you want your node to be able to run on multiple IBM Z® machines.

- To generate the Base64 encoded string, run the following command:

```
base64 <your-hostkey>.crt
```

Compared to guests not running IBM® Secure Execution, the first boot of the machine is longer because the entire image is encrypted with a randomly generated LUKS passphrase before the Ignition phase.

- Add Ignition protection

To protect the secrets that are stored in the Ignition config file from being read or even modified, you must encrypt the Ignition config file.

**NOTE**

To achieve the desired security, Ignition logging and local login are disabled by default when running IBM® Secure Execution.

- a. Fetch the public GPG key for the **secex-qemu.qcow2** image and encrypt the Ignition config with the key by running the following command:

```
gpg --recipient-file /path/to/ignition.gpg.pub --yes --output /path/to/config.ign.gpg --
verbose --armor --encrypt /path/to/config.ign
```

5. Follow the fast-track installation of RHCOS to install nodes by using the IBM® Secure Execution QCOW image.



NOTE

Before you start the VM, replace **serial=ignition** with **serial=ignition_crypted**, and add the **launchSecurity** parameter.

Verification

When you have completed the fast-track installation of RHCOS and Ignition runs at the first boot, verify if decryption is successful.

- If the decryption is successful, you can expect an output similar to the following example:

Example output

```
[ 2.801433] systemd[1]: Starting coreos-ignition-setup-user.service - CoreOS Ignition User
Config Setup...

[ 2.803959] coreos-secex-ignition-decrypt[731]: gpg: key <key_name>: public key "Secure
Execution (secex) 38.20230323.dev.0" imported
[ 2.808874] coreos-secex-ignition-decrypt[740]: gpg: encrypted with rsa4096 key, ID
<key_name>, created <yyyy-mm-dd>
[ OK ] Finished coreos-secex-igni...S Secex Ignition Config Decryptor.
```

- If the decryption fails, you can expect an output similar to the following example:

Example output

```
Starting coreos-ignition-s...reOS Ignition User Config Setup...
[ 2.863675] coreos-secex-ignition-decrypt[729]: gpg: key <key_name>: public key "Secure
Execution (secex) 38.20230323.dev.0" imported
[ 2.869178] coreos-secex-ignition-decrypt[738]: gpg: encrypted with RSA key, ID
<key_name>
[ 2.870347] coreos-secex-ignition-decrypt[738]: gpg: public key decryption failed: No secret
key
[ 2.870371] coreos-secex-ignition-decrypt[738]: gpg: decryption failed: No secret key
```

Additional resources

- [Introducing IBM® Secure Execution for Linux](#)
- [Linux as an IBM® Secure Execution host or guest](#)
- [Setting up IBM® Secure Execution on IBM Z](#)

18.5.11.2. Configuring NBDE with static IP in an IBM Z or IBM LinuxONE environment

Enabling NBDE disk encryption in an IBM Z® or IBM® LinuxONE environment requires additional steps, which are described in detail in this section.

Prerequisites

- You have set up the External Tang Server. See [Network-bound disk encryption](#) for instructions.
- You have installed the **butane** utility.
- You have reviewed the instructions for how to create machine configs with Butane.

Procedure

1. Create Butane configuration files for the control plane and compute nodes.
The following example of a Butane configuration for a control plane node creates a file named **master-storage.bu** for disk encryption:

```
variant: openshift
version: 4.16.0
metadata:
  name: master-storage
  labels:
    machineconfiguration.openshift.io/role: master
storage:
  luks:
    - clevis:
      tang:
        - thumbprint: QcPr_NHFJammnRCA3fFMVdNBwjs
          url: http://clevis.example.com:7500
      options: ①
        - --cipher
          - aes-cbc-essiv:sha256
    device: /dev/disk/by-partlabel/root
    label: luks-root
    name: root
    wipe_volume: true
  filesystems:
    - device: /dev/mapper/root
      format: xfs
      label: root
      wipe_filesystem: true
openshift:
  fips: true ②
```

- ① The cipher option is only required if FIPS mode is enabled. Omit the entry if FIPS is disabled.
- ② Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.

2. Create a customized initramfs file to boot the machine, by running the following command:

```
$ coreos-installer pxe customize \
  /root/rhcos-bootfiles/rhcos-<release>-live-initramfs.s390x.img \
  --dest-device /dev/disk/by-id/scsi-<serial_number> --dest-karg-append \
  ip=<ip_address>::<gateway_ip>:<subnet_mask>::<network_device>:none \
  --dest-karg-append nameserver=<nameserver_ip> \
  --dest-karg-append rd.neednet=1 -o \
  /root/rhcos-bootfiles/<node_name>-initramfs.s390x.img
```


**NOTE**

Before first boot, you must customize the initramfs for each node in the cluster, and add PXE kernel parameters.

3. Create a parameter file that includes **ignition.platform.id=metal** and **ignition.firstboot**.

```
rd.neednet=1 \
console=ttysclp0 \
ignition.firstboot ignition.platform.id=metal \
coreos.live.rootfs_url=http://<http_server>/rhcos-<version>-live-rootfs.<architecture>.img \ 1
coreos.inst.ignition_url=http://<http_server>/master.ign \ 2
ip=10.19.17.2::10.19.17.1:255.255.255.0::enb0d0:none nameserver=10.19.17.1 \
zfcplib.allow_lun_scan=0 \
rd.znet=qeth,0.0.bdd0,0.0.bdd1,0.0.bdd2,layer2=1 \
rd.zfcp=0.0.5677,0x600606680g7f0056,0x034F000000000000
```

- 1 Specify the location of the **rootfs** artifact for the **kernel** and **initramfs** you are booting. Only HTTP and HTTPS protocols are supported.
- 2 Specify the location of the Ignition config file. Use **master.ign** or **worker.ign**. Only HTTP and HTTPS protocols are supported.

**NOTE**

Write all options in the parameter file as a single line and make sure you have no newline characters.

Additional resources

- [Creating machine configs with Butane](#)

18.5.11.3. Fast-track installation by using a prepackaged QCOW2 disk image

Complete the following steps to create the machines in a fast-track installation of Red Hat Enterprise Linux CoreOS (RHCOS), importing a prepackaged Red Hat Enterprise Linux CoreOS (RHCOS) QEMU copy-on-write (QCOW2) disk image.

Prerequisites

- At least one LPAR running on RHEL 8.6 or later with KVM, referred to as RHEL KVM host in this procedure.
- The KVM/QEMU hypervisor is installed on the RHEL KVM host.

- A domain name server (DNS) that can perform hostname and reverse lookup for the nodes.
- A DHCP server that provides IP addresses.

Procedure

1. Obtain the RHEL QEMU copy-on-write (QCOW2) disk image file from the [Product Downloads](#) page on the Red Hat Customer Portal or from the [RHCOS image mirror](#) page.



IMPORTANT

The RHCOS images might not change with every release of OpenShift Container Platform. You must download images with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Only use the appropriate RHCOS QCOW2 image described in the following procedure.

2. Download the QCOW2 disk image and Ignition files to a common directory on the RHEL KVM host.

For example: `/var/lib/libvirt/images`



NOTE

The Ignition files are generated by the OpenShift Container Platform installer.

3. Create a new disk image with the QCOW2 disk image backing file for each KVM guest node.

```
$ qemu-img create -f qcow2 -F qcow2 -b /var/lib/libvirt/images/{source_rhcos_qemu}
/var/lib/libvirt/images/{vmname}.qcow2 {size}
```

4. Create the new KVM guest nodes using the Ignition file and the new disk image.

```
$ virt-install --noautoconsole \
  --connect qemu:///system \
  --name {vm_name} \
  --memory {memory} \
  --vcpus {vcpus} \
  --disk {disk} \
  --launchSecurity type="s390-pv" \ 1
  --import \
  --network network={network},mac={mac} \
  --disk path={ign_file},format=raw,readonly=on,serial=ignition,startup_policy=optional 2
```

1 If IBM® Secure Execution is enabled, add the **launchSecurity type="s390-pv"** parameter.

2 If IBM® Secure Execution is enabled, replace **serial=ignition** with **serial=ignition_crypted**.

18.5.11.4. Full installation on a new QCOW2 disk image

Complete the following steps to create the machines in a full installation on a new QEMU copy-on-write (QCOW2) disk image.

Prerequisites

- At least one LPAR running on RHEL 8.6 or later with KVM, referred to as RHEL KVM host in this procedure.
- The KVM/QEMU hypervisor is installed on the RHEL KVM host.
- A domain name server (DNS) that can perform hostname and reverse lookup for the nodes.
- An HTTP or HTTPS server is set up.

Procedure

1. Obtain the RHEL kernel, initramfs, and rootfs files from the [Product Downloads](#) page on the Red Hat Customer Portal or from the [RHCOS image mirror](#) page.



IMPORTANT

The RHCOS images might not change with every release of OpenShift Container Platform. You must download images with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Only use the appropriate RHCOS QCOW2 image described in the following procedure.

The file names contain the OpenShift Container Platform version number. They resemble the following examples:

- kernel: **rhcos-<version>-live-kernel-<architecture>**
 - initramfs: **rhcos-<version>-live-initramfs.<architecture>.img**
 - rootfs: **rhcos-<version>-live-rootfs.<architecture>.img**
2. Move the downloaded RHEL live kernel, initramfs, and rootfs as well as the Ignition files to an HTTP or HTTPS server before you launch **virt-install**.



NOTE

The Ignition files are generated by the OpenShift Container Platform installer.

3. Create the new KVM guest nodes using the RHEL kernel, initramfs, and Ignition files, the new disk image, and adjusted parm line arguments.
 - For **--location**, specify the location of the kernel/initrd on the HTTP or HTTPS server.
 - For **coreos.inst.ignition_url=**, specify the Ignition file for the machine role. Use **bootstrap.ign**, **master.ign**, or **worker.ign**. Only HTTP and HTTPS protocols are supported.
 - For **coreos.live.rootfs_url=**, specify the matching rootfs artifact for the kernel and initramfs you are booting. Only HTTP and HTTPS protocols are supported.

```
$ virt-install \
  --connect qemu:///system \
  --name {vm_name} \
  --vcpus {vcpus} \
  --memory {memory_mb} \
  --disk {vm_name}.qcow2,size={image_size| default(10,true)} \
  --network network={virt_network_parm} \
```

```

--boot hd \
--location {media_location},kernel={rhcos_kernel},initrd={rhcos_initrd} \
--extra-args "rd.neednet=1 coreos.inst.install_dev=/dev/vda coreos.live.rootfs_url=
{rhcos_liveos} ip={ip}::{default_gateway}:{subnet_mask_length}:{vm_name}:enc1:none:
{MTU} nameserver={dns} coreos.inst.ignition_url={rhcos_ign}" \
--noautoconsole \
--wait


```

18.5.11.5. Advanced RHCOS installation reference

This section illustrates the networking configuration and other advanced options that allow you to modify the Red Hat Enterprise Linux CoreOS (RHCOS) manual installation process. The following tables describe the kernel arguments and command-line options you can use with the RHCOS live installer and the **coreos-installer** command.

18.5.11.5.1. Networking options for ISO installations

If you install RHCOS from an ISO image, you can add kernel arguments manually when you boot the image to configure networking for a node. If no networking arguments are specified, DHCP is activated in the initramfs when RHCOS detects that networking is required to fetch the Ignition config file.



IMPORTANT

When adding networking arguments manually, you must also add the **rd.neednet=1** kernel argument to bring the network up in the initramfs.

The following information provides examples for configuring networking on your RHCOS nodes for ISO installations. The examples describe how to use the **ip=** and **nameserver=** kernel arguments.



NOTE

Ordering is important when adding the kernel arguments: **ip=** and **nameserver=**.

The networking options are passed to the **dracut** tool during system boot. For more information about the networking options supported by **dracut**, see the **dracut.cmdline** manual page.

The following examples are the networking options for ISO installation.

Configuring DHCP or static IP addresses

To configure an IP address, either use DHCP (**ip=dhcp**) or set an individual static IP address (**ip=<host_ip>**). If setting a static IP, you must then identify the DNS server IP address (**nameserver=<dns_ip>**) on each node. The following example sets:

- The node's IP address to **10.10.10.2**
- The gateway address to **10.10.10.254**
- The netmask to **255.255.255.0**
- The hostname to **core0.example.com**
- The DNS server address to **4.4.4.41**

- The auto-configuration value to **none**. No auto-configuration is required when IP networking is configured statically.

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none
nameserver=4.4.4.41
```



NOTE

When you use DHCP to configure IP addressing for the RHCOS machines, the machines also obtain the DNS server information through DHCP. For DHCP-based deployments, you can define the DNS server address that is used by the RHCOS nodes through your DHCP server configuration.

Configuring an IP address without a static hostname

You can configure an IP address without assigning a static hostname. If a static hostname is not set by the user, it will be picked up and automatically set by a reverse DNS lookup. To configure an IP address without a static hostname refer to the following example:

- The node's IP address to **10.10.10.2**
- The gateway address to **10.10.10.254**
- The netmask to **255.255.255.0**
- The DNS server address to **4.4.4.41**
- The auto-configuration value to **none**. No auto-configuration is required when IP networking is configured statically.

```
ip=10.10.10.2::10.10.10.254:255.255.255.0::enp1s0:none
nameserver=4.4.4.41
```

Specifying multiple network interfaces

You can specify multiple network interfaces by setting multiple **ip=** entries.

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none
ip=10.10.10.3::10.10.10.254:255.255.255.0:core0.example.com:enp2s0:none
```

Configuring default gateway and route

Optional: You can configure routes to additional networks by setting an **rd.route=** value.



NOTE

When you configure one or multiple networks, one default gateway is required. If the additional network gateway is different from the primary network gateway, the default gateway must be the primary network gateway.

- Run the following command to configure the default gateway:

```
ip=::10.10.10.254:::
```

- Enter the following command to configure the route for the additional network:


```
rd.route=20.20.20.0/24:20.20.20.254:enp2s0
```

Disabling DHCP on a single interface

You can disable DHCP on a single interface, such as when there are two or more network interfaces and only one interface is being used. In the example, the **enp1s0** interface has a static networking configuration and DHCP is disabled for **enp2s0**, which is not used:

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none
ip=::::core0.example.com:enp2s0:none
```

Combining DHCP and static IP configurations

You can combine DHCP and static IP configurations on systems with multiple network interfaces, for example:

```
ip=enp1s0:dhcp
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp2s0:none
```

Configuring VLANs on individual interfaces

Optional: You can configure VLANs on individual interfaces by using the **vlan=** parameter.

- To configure a VLAN on a network interface and use a static IP address, run the following command:

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp2s0.100:none
vlan=enp2s0.100:enp2s0
```

- To configure a VLAN on a network interface and to use DHCP, run the following command:

```
ip=enp2s0.100:dhcp
vlan=enp2s0.100:enp2s0
```

Providing multiple DNS servers

You can provide multiple DNS servers by adding a **nameserver=** entry for each server, for example:

```
nameserver=1.1.1.1
nameserver=8.8.8.8
```

18.5.12. Waiting for the bootstrap process to complete

The OpenShift Container Platform bootstrap process begins after the cluster nodes first boot into the persistent RHCOS environment that has been installed to disk. The configuration information provided through the Ignition config files is used to initialize the bootstrap process and install OpenShift Container Platform on the machines. You must wait for the bootstrap process to complete.

Prerequisites

- You have created the Ignition config files for your cluster.
- You have configured suitable network, DNS and load balancing infrastructure.
- You have obtained the installation program and generated the Ignition config files for your cluster.

- You installed RHCOS on your cluster machines and provided the Ignition config files that the OpenShift Container Platform installation program generated.

Procedure

- Monitor the bootstrap process:

```
$ ./openshift-install --dir <installation_directory> wait-for bootstrap-complete \ 1
--log-level=info 2
```

- For **<installation_directory>**, specify the path to the directory that you stored the installation files in.


- To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Example output

```
INFO Waiting up to 30m0s for the Kubernetes API at https://api.test.example.com:6443...
INFO API v1.29.4 up
INFO Waiting up to 30m0s for bootstrapping to complete...
INFO It is now safe to remove the bootstrap resources
```

The command succeeds when the Kubernetes API server signals that it has been bootstrapped on the control plane machines.

- After the bootstrap process is complete, remove the bootstrap machine from the load balancer.



IMPORTANT

You must remove the bootstrap machine from the load balancer at this point. You can also remove or reformat the bootstrap machine itself.

18.5.13. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

- Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

18.5.14. Approving the certificate signing requests for your machines

When you add machines to a cluster, two pending certificate signing requests (CSRs) are generated for each machine that you added. You must confirm that these CSRs are approved or, if necessary, approve them yourself. The client requests must be approved first, followed by the server requests.

Prerequisites

- You added machines to your cluster.

Procedure

1. Confirm that the cluster recognizes the machines:

```
$ oc get nodes
```

Example output

```
NAME      STATUS   ROLES    AGE   VERSION
master-0  Ready    master   63m   v1.29.4
master-1  Ready    master   63m   v1.29.4
master-2  Ready    master   64m   v1.29.4
```

The output lists all of the machines that you created.



NOTE

The preceding output might not include the compute nodes, also known as worker nodes, until some CSRs are approved.

2. Review the pending CSRs and ensure that you see the client requests with the **Pending** or **Approved** status for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

```
NAME      AGE   REQUESTOR                                     CONDITION
csr-8b2br  15m   system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper  Pending
csr-8vnps  15m   system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper  Pending
...
```

In this example, two machines are joining the cluster. You might see more approved CSRs in the list.

- If the CSRs were not approved, after all of the pending CSRs for the machines you added are in **Pending** status, approve the CSRs for your cluster machines:



NOTE

Because the CSRs rotate automatically, approve your CSRs within an hour of adding the machines to the cluster. If you do not approve them within an hour, the certificates will rotate, and more than two certificates will be present for each node. You must approve all of these certificates. After the client CSR is approved, the Kubelet creates a secondary CSR for the serving certificate, which requires manual approval. Then, subsequent serving certificate renewal requests are automatically approved by the **machine-approver** if the Kubelet requests a new certificate with identical parameters.



NOTE

For clusters running on platforms that are not machine API enabled, such as bare metal and other user-provisioned infrastructure, you must implement a method of automatically approving the kubelet serving certificate requests (CSRs). If a request is not approved, then the **oc exec**, **oc rsh**, and **oc logs** commands cannot succeed, because a serving certificate is required when the API server connects to the kubelet. Any operation that contacts the Kubelet endpoint requires this certificate approval to be in place. The method must watch for new CSRs, confirm that the CSR was submitted by the **node-bootstrap** service account in the **system:node** or **system:admin** groups, and confirm the identity of the node.

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{"\n"}\n{{end}}{{end}}' | xargs --no-run-if-empty oc adm certificate approve
```



NOTE

Some Operators might not become available until some CSRs are approved.

- Now that your client requests are approved, you must review the server requests for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

```

NAME      AGE   REQUESTOR                                CONDITION
csr-bfd72 5m26s system:node:ip-10-0-50-126.us-east-2.compute.internal
Pending
csr-c57lv 5m26s system:node:ip-10-0-95-157.us-east-2.compute.internal
Pending
...

```

5. If the remaining CSRs are not approved, and are in the **Pending** status, approve the CSRs for your cluster machines:

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

1 **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{"\n"}\n{{end}}' | xargs oc adm certificate approve
```

6. After all client and server CSRs have been approved, the machines have the **Ready** status. Verify this by running the following command:

```
$ oc get nodes
```

Example output

```

NAME      STATUS   ROLES    AGE   VERSION
master-0  Ready    master   73m   v1.29.4
master-1  Ready    master   73m   v1.29.4
master-2  Ready    master   74m   v1.29.4
worker-0  Ready    worker   11m   v1.29.4
worker-1  Ready    worker   11m   v1.29.4

```



NOTE

It can take a few minutes after approval of the server CSRs for the machines to transition to the **Ready** status.

Additional information

- For more information on CSRs, see [Certificate Signing Requests](#).

18.5.15. Initial Operator configuration

After the control plane initializes, you must immediately configure some Operators so that they all become available.

Prerequisites

- Your control plane has initialized.

Procedure

1. Watch the cluster components come online:

```
$ watch -n5 oc get clusteroperators
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.16.0	True	False	False	19m
baremetal	4.16.0	True	False	False	37m
cloud-credential	4.16.0	True	False	False	40m
cluster-autoscaler	4.16.0	True	False	False	37m
config-operator	4.16.0	True	False	False	38m
console	4.16.0	True	False	False	26m
csi-snapshot-controller	4.16.0	True	False	False	37m
dns	4.16.0	True	False	False	37m
etcd	4.16.0	True	False	False	36m
image-registry	4.16.0	True	False	False	31m
ingress	4.16.0	True	False	False	30m
insights	4.16.0	True	False	False	31m
kube-apiserver	4.16.0	True	False	False	26m
kube-controller-manager	4.16.0	True	False	False	36m
kube-scheduler	4.16.0	True	False	False	36m
kube-storage-version-migrator	4.16.0	True	False	False	37m
machine-api	4.16.0	True	False	False	29m
machine-approver	4.16.0	True	False	False	37m
machine-config	4.16.0	True	False	False	36m
marketplace	4.16.0	True	False	False	37m
monitoring	4.16.0	True	False	False	29m
network	4.16.0	True	False	False	38m
node-tuning	4.16.0	True	False	False	37m
openshift-apiserver	4.16.0	True	False	False	32m
openshift-controller-manager	4.16.0	True	False	False	30m
openshift-samples	4.16.0	True	False	False	32m
operator-lifecycle-manager	4.16.0	True	False	False	37m
operator-lifecycle-manager-catalog	4.16.0	True	False	False	37m
operator-lifecycle-manager-packageserver	4.16.0	True	False	False	32m
service-ca	4.16.0	True	False	False	38m
storage	4.16.0	True	False	False	37m

2. Configure the Operators that are not available.

18.5.15.1. Disabling the default OperatorHub catalog sources

Operator catalogs that source content provided by Red Hat and community projects are configured for OperatorHub by default during an OpenShift Container Platform installation. In a restricted network environment, you must disable the default catalogs as a cluster administrator.

Procedure

- Disable the sources for the default catalogs by adding **disableAllDefaultSources: true** to the **OperatorHub** object:

```
$ oc patch OperatorHub cluster --type json \
  -p '[{"op": "add", "path": "/spec/disableAllDefaultSources", "value": true}]'
```

TIP

Alternatively, you can use the web console to manage catalog sources. From the **Administration** → **Cluster Settings** → **Configuration** → **OperatorHub** page, click the **Sources** tab, where you can create, update, delete, disable, and enable individual sources.

18.5.15.2. Image registry storage configuration

The Image Registry Operator is not initially available for platforms that do not provide default storage. After installation, you must configure your registry to use storage so that the Registry Operator is made available.

Instructions are shown for configuring a persistent volume, which is required for production clusters. Where applicable, instructions are shown for configuring an empty directory as the storage location, which is available for only non-production clusters.

Additional instructions are provided for allowing the image registry to use block storage types by using the **Recreate** rollout strategy during upgrades.

18.5.15.2.1. Configuring registry storage for IBM Z

As a cluster administrator, following installation you must configure your registry to use storage.

Prerequisites

- You have access to the cluster as a user with the **cluster-admin** role.
- You have a cluster on IBM Z®.
- You have provisioned persistent storage for your cluster, such as Red Hat OpenShift Data Foundation.




IMPORTANT

OpenShift Container Platform supports **ReadWriteOnce** access for image registry storage when you have only one replica. **ReadWriteOnce** access also requires that the registry uses the **Recreate** rollout strategy. To deploy an image registry that supports high availability with two or more replicas, **ReadWriteMany** access is required.

- Must have 100Gi capacity.

Procedure

1. To configure your registry to use storage, change the **spec.storage.pvc** in the **configs.imageregistry/cluster** resource.

**NOTE**


When you use shared storage, review your security settings to prevent outside access.

- Verify that you do not have a registry pod:

```
$ oc get pod -n openshift-image-registry -l docker-registry=default
```

Example output

```
No resources found in openshift-image-registry namespace
```

**NOTE**

If you do have a registry pod in your output, you do not need to continue with this procedure.

- Check the registry configuration:

```
$ oc edit configs.imageregistry.operator.openshift.io
```

Example output

```
storage:
  pvc:
    claim:
```

Leave the **claim** field blank to allow the automatic creation of an **image-registry-storage** PVC.

- Check the **clusteroperator** status:

```
$ oc get clusteroperator image-registry
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
image-registry	4.16	True	False	False	6h50m

- Ensure that your registry is set to managed to enable building and pushing of images.

- Run:

```
$ oc edit configs.imageregistry/cluster
```

Then, change the line

```
managementState: Removed
```

to

managementState: Managed

18.5.15.2.2. Configuring storage for the image registry in non-production clusters

You must configure storage for the Image Registry Operator. For non-production clusters, you can set the image registry to an empty directory. If you do so, all images are lost if you restart the registry.

Procedure

- To set the image registry storage to an empty directory:

```
$ oc patch configs.imageregistry.operator.openshift.io cluster --type merge --patch '{"spec": {"storage":{"emptyDir":{}}}'
```



WARNING

Configure this option for only non-production clusters.

If you run this command before the Image Registry Operator initializes its components, the **oc patch** command fails with the following error:

```
Error from server (NotFound): configs.imageregistry.operator.openshift.io "cluster" not found
```

Wait a few minutes and run the command again.

18.5.16. Completing installation on user-provisioned infrastructure

After you complete the Operator configuration, you can finish installing the cluster on infrastructure that you provide.

Prerequisites

- Your control plane has initialized.
- You have completed the initial Operator configuration.

Procedure

- Confirm that all the cluster components are online with the following command:

```
$ watch -n5 oc get clusteroperators
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED
authentication	4.16.0	True	False	False 19m

baremetal	4.16.0	True	False	False	37m
cloud-credential	4.16.0	True	False	False	40m
cluster-autoscaler	4.16.0	True	False	False	37m
config-operator	4.16.0	True	False	False	38m
console	4.16.0	True	False	False	26m
csi-snapshot-controller	4.16.0	True	False	False	37m
dns	4.16.0	True	False	False	37m
etcd	4.16.0	True	False	False	36m
image-registry	4.16.0	True	False	False	31m
ingress	4.16.0	True	False	False	30m
insights	4.16.0	True	False	False	31m
kube-apiserver	4.16.0	True	False	False	26m
kube-controller-manager	4.16.0	True	False	False	36m
kube-scheduler	4.16.0	True	False	False	36m
kube-storage-version-migrator	4.16.0	True	False	False	37m
machine-api	4.16.0	True	False	False	29m
machine-approver	4.16.0	True	False	False	37m
machine-config	4.16.0	True	False	False	36m
marketplace	4.16.0	True	False	False	37m
monitoring	4.16.0	True	False	False	29m
network	4.16.0	True	False	False	38m
node-tuning	4.16.0	True	False	False	37m
openshift-apiserver	4.16.0	True	False	False	32m
openshift-controller-manager	4.16.0	True	False	False	30m
openshift-samples	4.16.0	True	False	False	32m
operator-lifecycle-manager	4.16.0	True	False	False	37m
operator-lifecycle-manager-catalog	4.16.0	True	False	False	37m
operator-lifecycle-manager-packageserver	4.16.0	True	False	False	32m
service-ca	4.16.0	True	False	False	38m
storage	4.16.0	True	False	False	37m

Alternatively, the following command notifies you when all of the clusters are available. It also retrieves and displays credentials:

```
$ ./openshift-install --dir <installation_directory> wait-for install-complete 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

Example output

```
INFO Waiting up to 30m0s for the cluster to initialize...
```

The command succeeds when the Cluster Version Operator finishes deploying the OpenShift Container Platform cluster from Kubernetes API server.



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

2. Confirm that the Kubernetes API server is communicating with the pods.

a. To view a list of all pods, use the following command:

```
$ oc get pods --all-namespaces
```

Example output

```

NAMESPACE           NAME                                     READY  STATUS
RESTARTS  AGE
openshift-apiserver-operator  openshift-apiserver-operator-85cb746d55-zqhs8  1/1
Running   1    9m
openshift-apiserver          apiserver-67b9g                                1/1  Running  0
3m
openshift-apiserver          apiserver-ljcmx                                1/1  Running  0
1m
openshift-apiserver          apiserver-z25h4                                1/1  Running  0
2m
openshift-authentication-operator  authentication-operator-69d5d8bf84-vh2n8  1/1
Running   0    5m
...

```

b. View the logs for a pod that is listed in the output of the previous command by using the following command:

```
$ oc logs <pod_name> -n <namespace> ❶
```

- ❶ Specify the pod name and namespace, as shown in the output of the previous command.

If the pod logs display, the Kubernetes API server can communicate with the cluster machines.

3. For an installation with Fibre Channel Protocol (FCP), additional steps are required to enable multipathing. Do not enable multipathing during installation. See "Enabling multipathing with kernel arguments on RHCOS" in the *Postinstallation machine configuration tasks* documentation for more information.

4. Register your cluster on the [Cluster registration](#) page.

Additional resources

- [How to generate SOSREPORT within OpenShift Container Platform version 4 nodes without SSH.](#)

18.5.17. Next steps

- [Customize your cluster.](#)
- If the mirror registry that you used to install your cluster has a trusted CA, add it to the cluster by [configuring additional trust stores](#).
- If necessary, you can [opt out of remote health reporting](#) .
- If necessary, see [Registering your disconnected cluster](#)


18.6. INSTALLING A CLUSTER IN AN LPAR ON IBM Z AND IBM LINUXONE

In OpenShift Container Platform version 4.16, you can install a cluster in a logical partition (LPAR) on IBM Z® or IBM® LinuxONE infrastructure that you provision.



NOTE

While this document refers only to IBM Z®, all information in it also applies to IBM® LinuxONE.




IMPORTANT

Additional considerations exist for non-bare metal platforms. Review the information in the [guidelines for deploying OpenShift Container Platform on non-tested platforms](#) before you install an OpenShift Container Platform cluster.

18.6.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- Before you begin the installation process, you must clean the installation directory. This ensures that the required installation files are created and updated during the installation process.
- You provisioned [persistent storage using OpenShift Data Foundation](#) or other supported storage protocols for your cluster. To deploy a private image registry, you must set up persistent storage with **ReadWriteMany** access.
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.

**NOTE**


Be sure to also review this site list if you are configuring a proxy.

18.6.2. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.

**IMPORTANT**

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

18.6.3. Requirements for a cluster with user-provisioned infrastructure

For a cluster that contains user-provisioned infrastructure, you must deploy all of the required machines.

This section describes the requirements for deploying OpenShift Container Platform on user-provisioned infrastructure.

18.6.3.1. Required machines for cluster installation

The smallest OpenShift Container Platform clusters require the following hosts:

Table 18.76. Minimum required hosts

Hosts	Description
One temporary bootstrap machine	The cluster requires the bootstrap machine to deploy the OpenShift Container Platform cluster on the three control plane machines. You can remove the bootstrap machine after you install the cluster.
Three control plane machines	The control plane machines run the Kubernetes and OpenShift Container Platform services that form the control plane.
At least two compute machines, which are also known as worker machines.	The workloads requested by OpenShift Container Platform users run on the compute machines.

**IMPORTANT**

To maintain high availability of your cluster, use separate physical hosts for these cluster machines.

The bootstrap and control plane machines must use Red Hat Enterprise Linux CoreOS (RHCOS) as the operating system. However, the compute machines can choose between Red Hat Enterprise Linux CoreOS (RHCOS), Red Hat Enterprise Linux (RHEL) 8.6 and later.

Note that RHCOS is based on Red Hat Enterprise Linux (RHEL) 9.2 and inherits all of its hardware certifications and requirements. See [Red Hat Enterprise Linux technology capabilities and limits](#).

18.6.3.2. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 18.77. Minimum resource requirements

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)
Bootstrap	RHCOS	4	16 GB	100 GB	N/A
Control plane	RHCOS	4	16 GB	100 GB	N/A
Compute	RHCOS	2	8 GB	100 GB	N/A

1. One physical core (IFL) provides two logical cores (threads) when SMT-2 is enabled. The hypervisor can provide two or more vCPUs.

**NOTE**

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

18.6.3.3. Minimum IBM Z system environment

You can install OpenShift Container Platform version 4.16 on the following IBM® hardware:

- IBM® z16 (all models), IBM® z15 (all models), IBM® z14 (all models)
- IBM® LinuxONE 4 (all models), IBM® LinuxONE III (all models), IBM® LinuxONE Emperor II, IBM® LinuxONE Rockhopper II



IMPORTANT

When running OpenShift Container Platform on IBM Z® without a hypervisor use the Dynamic Partition Manager (DPM) to manage your machine.

Hardware requirements

- The equivalent of six Integrated Facilities for Linux (IFL), which are SMT2 enabled, for each cluster.
- At least one network connection to both connect to the **LoadBalancer** service and to serve data for traffic outside the cluster.



NOTE

You can use dedicated or shared IFLs to assign sufficient compute resources. Resource sharing is one of the key strengths of IBM Z®. However, you must adjust capacity correctly on each hypervisor layer and ensure sufficient resources for every OpenShift Container Platform cluster.



IMPORTANT

Since the overall performance of the cluster can be impacted, the LPARs that are used to set up the OpenShift Container Platform clusters must provide sufficient compute capacity. In this context, LPAR weight management, entitlements, and CPU shares on the hypervisor level play an important role.

Operating system requirements

- Five logical partitions (LPARs)
 - Three LPARs for OpenShift Container Platform control plane machines
 - Two LPARs for OpenShift Container Platform compute machines
- One machine for the temporary OpenShift Container Platform bootstrap machine

IBM Z network connectivity requirements

To install on IBM Z® in an LPAR, you need:

- A direct-attached OSA or RoCE network adapter
- For a preferred setup, use OSA link aggregation.

Disk storage

- FICON attached disk storage (DASDs). These can be dedicated DASDs that must be formatted

as CDL, which is the default. To reach the minimum required DASD size for Red Hat Enterprise Linux CoreOS (RHCOS) installations, you need extended address volumes (EAV). If available, use HyperPAV to ensure optimal performance.

- FCP attached disk storage

Storage / Main Memory

- 16 GB for OpenShift Container Platform control plane machines
- 8 GB for OpenShift Container Platform compute machines
- 16 GB for the temporary OpenShift Container Platform bootstrap machine

Additional resources

- [Processors Resource/Systems Manager Planning Guide](#) in IBM® Documentation for PR/SM mode considerations.
- [IBM Dynamic Partition Manager \(DPM\) Guide](#) in IBM® Documentation for DPM mode considerations.
- [Topics in LPAR performance](#) for LPAR weight management and entitlements.
- [Recommended host practices for IBM Z® & IBM® LinuxONE environments](#)

18.6.3.4. Preferred IBM Z system environment

Hardware requirements

- Three LPARS that each have the equivalent of six IFLs, which are SMT2 enabled, for each cluster.
- Two network connections to both connect to the **LoadBalancer** service and to serve data for traffic outside the cluster.
- HiperSockets that are attached to a node directly as a device. To directly connect HiperSockets to a node, you must set up a gateway to the external network via a RHEL 8 guest to bridge to the HiperSockets network.

Operating system requirements

- Three LPARs for OpenShift Container Platform control plane machines.
- At least six LPARs for OpenShift Container Platform compute machines.
- One machine or LPAR for the temporary OpenShift Container Platform bootstrap machine.

IBM Z network connectivity requirements

To install on IBM Z® in an LPAR, you need:

- A direct-attached OSA or RoCE network adapter
- For a preferred setup, use OSA link aggregation.

Disk storage

- FICON attached disk storage (DASDs). These can be dedicated DASDs that must be formatted

as CDL, which is the default. To reach the minimum required DASD size for Red Hat Enterprise Linux CoreOS (RHCOS) installations, you need extended address volumes (EAV). If available, use HyperPAV to ensure optimal performance.

- FCP attached disk storage

Storage / Main Memory

- 16 GB for OpenShift Container Platform control plane machines
- 8 GB for OpenShift Container Platform compute machines
- 16 GB for the temporary OpenShift Container Platform bootstrap machine

18.6.3.5. Certificate signing requests management

Because your cluster has limited access to automatic machine management when you use infrastructure that you provision, you must provide a mechanism for approving cluster certificate signing requests (CSRs) after installation. The **kube-controller-manager** only approves the kubelet client CSRs. The **machine-approver** cannot guarantee the validity of a serving certificate that is requested by using kubelet credentials because it cannot confirm that the correct machine issued the request. You must determine and implement a method of verifying the validity of the kubelet serving certificate requests and approving them.

18.6.3.6. Networking requirements for user-provisioned infrastructure

All the Red Hat Enterprise Linux CoreOS (RHCOS) machines require networking to be configured in **initramfs** during boot to fetch their Ignition config files.

During the initial boot, the machines require an HTTP or HTTPS server to establish a network connection to download their Ignition config files.

The machines are configured with static IP addresses. No DHCP server is required. Ensure that the machines have persistent IP addresses and hostnames.

The Kubernetes API server must be able to resolve the node names of the cluster machines. If the API servers and worker nodes are in different zones, you can configure a default DNS search zone to allow the API server to resolve the node names. Another supported approach is to always refer to hosts by their fully-qualified domain names in both the node objects and all DNS requests.

18.6.3.6.1. Network connectivity requirements

You must configure the network connectivity between machines to allow OpenShift Container Platform cluster components to communicate. Each machine must be able to resolve the hostnames of all other machines in the cluster.

This section provides details about the ports that are required.



IMPORTANT

In connected OpenShift Container Platform environments, all nodes are required to have internet access to pull images for platform containers and provide telemetry data to Red Hat.

Table 18.78. Ports used for all-machine to all-machine communications

Protocol	Port	Description
ICMP	N/A	Network reachability tests
TCP	1936	Metrics
	9000-9999	Host level services, including the node exporter on ports 9100-9101 and the Cluster Version Operator on port 9099 .
	10250-10259	The default ports that Kubernetes reserves
UDP	4789	VXLAN
	6081	Geneve
	9000-9999	Host level services, including the node exporter on ports 9100-9101 .
	500	IPsec IKE packets
	4500	IPsec NAT-T packets
	123	Network Time Protocol (NTP) on UDP port 123 If an external NTP time server is configured, you must open UDP port 123 .
TCP/UDP	30000-32767	Kubernetes node port
ESP	N/A	IPsec Encapsulating Security Payload (ESP)

Table 18.79. Ports used for all-machine to control plane communications

Protocol	Port	Description
TCP	6443	Kubernetes API

Table 18.80. Ports used for control plane machine to control plane machine communications

Protocol	Port	Description
TCP	2379-2380	etcd server and peer ports

NTP configuration for user-provisioned infrastructure

OpenShift Container Platform clusters are configured to use a public Network Time Protocol (NTP) server by default. If you want to use a local enterprise NTP server, or if your cluster is being deployed in a disconnected network, you can configure the cluster to use a specific time server. For more information,

see the documentation for *Configuring chrony time service*.

Additional resources

- [Configuring chrony time service](#)

18.6.3.7. User-provisioned DNS requirements

In OpenShift Container Platform deployments, DNS name resolution is required for the following components:


- The Kubernetes API
- The OpenShift Container Platform application wildcard
- The bootstrap, control plane, and compute machines

Reverse DNS resolution is also required for the Kubernetes API, the bootstrap machine, the control plane machines, and the compute machines.

DNS A/AAAA or CNAME records are used for name resolution and PTR records are used for reverse name resolution. The reverse records are important because Red Hat Enterprise Linux CoreOS (RHCOS) uses the reverse records to set the hostnames for all the nodes, unless the hostnames are provided by DHCP. Additionally, the reverse records are used to generate the certificate signing requests (CSR) that OpenShift Container Platform needs to operate.

The following DNS records are required for a user-provisioned OpenShift Container Platform cluster and they must be in place before installation. In each record, **<cluster_name>** is the cluster name and **<base_domain>** is the base domain that you specify in the **install-config.yaml** file. A complete DNS record takes the form: **<component>.<cluster_name>.<base_domain>..**

Table 18.81. Required DNS records

Component	Record	Description
Kubernetes API	api.<cluster_name>.<base_domain>..	A DNS A/AAAA or CNAME record, and a DNS PTR record, to identify the API load balancer. These records must be resolvable by both clients external to the cluster and from all the nodes within the cluster.
	api-int.<cluster_name>.<base_domain>..	A DNS A/AAAA or CNAME record, and a DNS PTR record, to internally identify the API load balancer. These records must be resolvable from all the nodes within the cluster.
		 <p>IMPORTANT</p> <p>The API server must be able to resolve the worker nodes by the hostnames that are recorded in Kubernetes. If the API server cannot resolve the node names, then proxied API calls can fail, and you cannot retrieve logs from pods.</p>

Component	Record	Description
Routes	*.apps.<cluster_name>.<base_domain>.	A wildcard DNS A/AAAA or CNAME record that refers to the application ingress load balancer. The application ingress load balancer targets the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default. These records must be resolvable by both clients external to the cluster and from all the nodes within the cluster. For example, console-openshift-console.apps.<cluster_name>.<base_domain> is used as a wildcard route to the OpenShift Container Platform console.
Bootstrap machine	bootstrap.<cluster_name>.<base_domain>.	A DNS A/AAAA or CNAME record, and a DNS PTR record, to identify the bootstrap machine. These records must be resolvable by the nodes within the cluster.
Control plane machines	<control_plane><n>.<cluster_name>.<base_domain>.	DNS A/AAAA or CNAME records and DNS PTR records to identify each machine for the control plane nodes. These records must be resolvable by the nodes within the cluster.
Compute machines	<compute><n>.<cluster_name>.<base_domain>.	DNS A/AAAA or CNAME records and DNS PTR records to identify each machine for the worker nodes. These records must be resolvable by the nodes within the cluster.

**NOTE**

In OpenShift Container Platform 4.4 and later, you do not need to specify etcd host and SRV records in your DNS configuration.

TIP

You can use the **dig** command to verify name and reverse name resolution. See the section on *Validating DNS resolution for user-provisioned infrastructure* for detailed validation steps.

18.6.3.7.1. Example DNS configuration for user-provisioned clusters

This section provides A and PTR record configuration samples that meet the DNS requirements for deploying OpenShift Container Platform on user-provisioned infrastructure. The samples are not meant to provide advice for choosing one DNS solution over another.

In the examples, the cluster name is **ocp4** and the base domain is **example.com**.

Example DNS A record configuration for a user-provisioned cluster

The following example is a BIND zone file that shows sample A records for name resolution in a user-provisioned cluster.

Example 18.13. Sample DNS zone database

```

$TTL 1W
@ IN SOA ns1.example.com. root (
    2019070700 ; serial
    3H ; refresh (3 hours)
    30M ; retry (30 minutes)
    2W ; expiry (2 weeks)
    1W ) ; minimum (1 week)
IN NS ns1.example.com.
IN MX 10 smtp.example.com.
;
;
;
ns1.example.com. IN A 192.168.1.5
smtp.example.com. IN A 192.168.1.5
;
;
helper.example.com. IN A 192.168.1.5
helper.ocp4.example.com. IN A 192.168.1.5
;
;
api.ocp4.example.com. IN A 192.168.1.5 1
api-int.ocp4.example.com. IN A 192.168.1.5 2
;
;
*.apps.ocp4.example.com. IN A 192.168.1.5 3
;
;
bootstrap.ocp4.example.com. IN A 192.168.1.96 4
;
;
control-plane0.ocp4.example.com. IN A 192.168.1.97 5
control-plane1.ocp4.example.com. IN A 192.168.1.98 6
control-plane2.ocp4.example.com. IN A 192.168.1.99 7
;
;
compute0.ocp4.example.com. IN A 192.168.1.11 8
compute1.ocp4.example.com. IN A 192.168.1.7 9
;
;
;EOF

```

- 1 Provides name resolution for the Kubernetes API. The record refers to the IP address of the API load balancer.
- 2 Provides name resolution for the Kubernetes API. The record refers to the IP address of the API load balancer and is used for internal cluster communications.
- 3 Provides name resolution for the wildcard routes. The record refers to the IP address of the application ingress load balancer. The application ingress load balancer targets the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.

**NOTE**

In the example, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

- 4 Provides name resolution for the bootstrap machine.
- 5 6 7 Provides name resolution for the control plane machines.
- 8 9 Provides name resolution for the compute machines.

Example DNS PTR record configuration for a user-provisioned cluster

The following example BIND zone file shows sample PTR records for reverse name resolution in a user-provisioned cluster.

Example 18.14. Sample DNS zone database for reverse records

```
$TTL 1W
@ IN SOA ns1.example.com. root (
    2019070700 ; serial
    3H ; refresh (3 hours)
    30M ; retry (30 minutes)
    2W ; expiry (2 weeks)
    1W ) ; minimum (1 week)
IN NS ns1.example.com.
;
5.1.168.192.in-addr.arpa. IN PTR api.ocp4.example.com. 1
5.1.168.192.in-addr.arpa. IN PTR api-int.ocp4.example.com. 2
;
96.1.168.192.in-addr.arpa. IN PTR bootstrap.ocp4.example.com. 3
;
97.1.168.192.in-addr.arpa. IN PTR control-plane0.ocp4.example.com. 4
98.1.168.192.in-addr.arpa. IN PTR control-plane1.ocp4.example.com. 5
99.1.168.192.in-addr.arpa. IN PTR control-plane2.ocp4.example.com. 6
;
11.1.168.192.in-addr.arpa. IN PTR compute0.ocp4.example.com. 7
7.1.168.192.in-addr.arpa. IN PTR compute1.ocp4.example.com. 8
;
;EOF
```

- 1 Provides reverse DNS resolution for the Kubernetes API. The PTR record refers to the record name of the API load balancer.
- 2 Provides reverse DNS resolution for the Kubernetes API. The PTR record refers to the record name of the API load balancer and is used for internal cluster communications.
- 3 Provides reverse DNS resolution for the bootstrap machine.
- 4 5 6 Provides reverse DNS resolution for the control plane machines.
- 7 8 Provides reverse DNS resolution for the compute machines.

**NOTE**

A PTR record is not required for the OpenShift Container Platform application wildcard.

18.6.3.8. Load balancing requirements for user-provisioned infrastructure

Before you install OpenShift Container Platform, you must provision the API and application Ingress load balancing infrastructure. In production scenarios, you can deploy the API and application Ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

**NOTE**

If you want to deploy the API and application Ingress load balancers with a Red Hat Enterprise Linux (RHEL) instance, you must purchase the RHEL subscription separately.

The load balancing infrastructure must meet the following requirements:

1. **API load balancer.** Provides a common endpoint for users, both human and machine, to interact with and configure the platform. Configure the following conditions:
 - Layer 4 load balancing only. This can be referred to as Raw TCP or SSL Passthrough mode.
 - A stateless load balancing algorithm. The options vary based on the load balancer implementation.

**IMPORTANT**

Do not configure session persistence for an API load balancer. Configuring session persistence for a Kubernetes API server might cause performance issues from excess application traffic for your OpenShift Container Platform cluster and the Kubernetes API that runs inside the cluster.

Configure the following ports on both the front and back of the load balancers:

Table 18.82. API load balancer

Port	Back-end machines (pool members)	Internal	External	Description
6443	Bootstrap and control plane. You remove the bootstrap machine from the load balancer after the bootstrap machine initializes the cluster control plane. You must configure the /readyz endpoint for the API server health check probe.	X	X	Kubernetes API server
22623	Bootstrap and control plane. You remove the bootstrap machine from the load balancer after the bootstrap machine initializes the cluster control plane.	X		Machine config server

**NOTE**

The load balancer must be configured to take a maximum of 30 seconds from the time the API server turns off the **/readyz** endpoint to the removal of the API server instance from the pool. Within the time frame after **/readyz** returns an error or becomes healthy, the endpoint must have been removed or added. Probing every 5 or 10 seconds, with two successful requests to become healthy and three to become unhealthy, are well-tested values.

2. **Application Ingress load balancer.** Provides an ingress point for application traffic flowing in from outside the cluster. A working configuration for the Ingress router is required for an OpenShift Container Platform cluster.

Configure the following conditions:

- Layer 4 load balancing only. This can be referred to as Raw TCP or SSL Passthrough mode.
- A connection-based or session-based persistence is recommended, based on the options available and types of applications that will be hosted on the platform.

TIP

If the true IP address of the client can be seen by the application Ingress load balancer, enabling source IP-based session persistence can improve performance for applications that use end-to-end TLS encryption.

Configure the following ports on both the front and back of the load balancers:

Table 18.83. Application Ingress load balancer

Port	Back-end machines (pool members)	Internal	External	Description
443	The machines that run the Ingress Controller pods, compute, or worker, by default.	X	X	HTTPS traffic
80	The machines that run the Ingress Controller pods, compute, or worker, by default.	X	X	HTTP traffic

**NOTE**

If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application Ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes.

18.6.3.8.1. Example load balancer configuration for user-provisioned clusters

This section provides an example API and application Ingress load balancer configuration that meets the load balancing requirements for user-provisioned clusters. The sample is an **/etc/haproxy/haproxy.cfg** configuration for an HAProxy load balancer. The example is not meant to provide advice for choosing one load balancing solution over another.

In the example, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.



NOTE

If you are using HAProxy as a load balancer and SELinux is set to **enforcing**, you must ensure that the HAProxy service can bind to the configured TCP port by running **setsebool -P haproxy_connect_any=1**.

Example 18.15. Sample API and application Ingress load balancer configuration

```

global
  log      127.0.0.1 local2
  pidfile  /var/run/haproxy.pid
  maxconn 4000
  daemon
defaults
  mode          http
  log           global
  option        dontlognull
  option http-server-close
  option        redispatch
  retries       3
  timeout http-request 10s
  timeout queue   1m
  timeout connect 10s
  timeout client  1m
  timeout server  1m
  timeout http-keep-alive 10s
  timeout check   10s
  maxconn        3000
listen api-server-6443 1
  bind *:6443
  mode tcp
  option httpchk GET /readyz HTTP/1.0
  option log-health-checks
  balance roundrobin
  server bootstrap bootstrap.ocp4.example.com:6443 verify none check check-ssl inter 10s fall 2
  rise 3 backup 2
  server master0 master0.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
  fall 2 rise 3
  server master1 master1.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
  fall 2 rise 3
  server master2 master2.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
  fall 2 rise 3
listen machine-config-server-22623 3
  bind *:22623
  mode tcp
  server bootstrap bootstrap.ocp4.example.com:22623 check inter 1s backup 4
  server master0 master0.ocp4.example.com:22623 check inter 1s
  server master1 master1.ocp4.example.com:22623 check inter 1s
  server master2 master2.ocp4.example.com:22623 check inter 1s
listen ingress-router-443 5


```

```

bind *:443
mode tcp
balance source
server compute0 compute0.ocp4.example.com:443 check inter 1s
server compute1 compute1.ocp4.example.com:443 check inter 1s
listen ingress-router-80 6
bind *:80
mode tcp
balance source
server compute0 compute0.ocp4.example.com:80 check inter 1s
server compute1 compute1.ocp4.example.com:80 check inter 1s

```

- 1** Port **6443** handles the Kubernetes API traffic and points to the control plane machines.
- 2** **4** The bootstrap entries must be in place before the OpenShift Container Platform cluster installation and they must be removed after the bootstrap process is complete.
- 3** Port **22623** handles the machine config server traffic and points to the control plane machines.
- 5** Port **443** handles the HTTPS traffic and points to the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.
- 6** Port **80** handles the HTTP traffic and points to the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.



NOTE

If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application Ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes.

TIP

If you are using HAProxy as a load balancer, you can check that the **haproxy** process is listening on ports **6443**, **22623**, **443**, and **80** by running **netstat -nltp** on the HAProxy node.

18.6.4. Preparing the user-provisioned infrastructure

Before you install OpenShift Container Platform on user-provisioned infrastructure, you must prepare the underlying infrastructure.

This section provides details about the high-level steps required to set up your cluster infrastructure in preparation for an OpenShift Container Platform installation. This includes configuring IP networking and network connectivity for your cluster nodes, preparing a web server for the Ignition files, enabling the required ports through your firewall, and setting up the required DNS and load balancing infrastructure.


After preparation, your cluster infrastructure must meet the requirements outlined in the *Requirements for a cluster with user-provisioned infrastructure* section.

Prerequisites

- You have reviewed the [OpenShift Container Platform 4.x Tested Integrations](#) page.
- You have reviewed the infrastructure requirements detailed in the *Requirements for a cluster with user-provisioned infrastructure* section.

Procedure

1. Set up static IP addresses.
2. Set up an HTTP or HTTPS server to provide Ignition files to the cluster nodes.
3. Ensure that your network infrastructure provides the required network connectivity between the cluster components. See the *Networking requirements for user-provisioned infrastructure* section for details about the requirements.
4. Configure your firewall to enable the ports required for the OpenShift Container Platform cluster components to communicate. See *Networking requirements for user-provisioned infrastructure* section for details about the ports that are required.



IMPORTANT

By default, port **1936** is accessible for an OpenShift Container Platform cluster, because each control plane node needs access to this port.

Avoid using the Ingress load balancer to expose this port, because doing so might result in the exposure of sensitive information, such as statistics and metrics, related to Ingress Controllers.


5. Setup the required DNS infrastructure for your cluster.
 - a. Configure DNS name resolution for the Kubernetes API, the application wildcard, the bootstrap machine, the control plane machines, and the compute machines.
 - b. Configure reverse DNS resolution for the Kubernetes API, the bootstrap machine, the control plane machines, and the compute machines.
See the *User-provisioned DNS requirements* section for more information about the OpenShift Container Platform DNS requirements.
6. Validate your DNS configuration.
 - a. From your installation node, run DNS lookups against the record names of the Kubernetes API, the wildcard routes, and the cluster nodes. Validate that the IP addresses in the responses correspond to the correct components.
 - b. From your installation node, run reverse DNS lookups against the IP addresses of the load balancer and the cluster nodes. Validate that the record names in the responses correspond to the correct components.
See the *Validating DNS resolution for user-provisioned infrastructure* section for detailed DNS validation steps.
7. Provision the required API and application ingress load balancing infrastructure. See the *Load balancing requirements for user-provisioned infrastructure* section for more information about the requirements.

**NOTE**

Some load balancing solutions require the DNS name resolution for the cluster nodes to be in place before the load balancing is initialized.

18.6.5. Validating DNS resolution for user-provisioned infrastructure

You can validate your DNS configuration before installing OpenShift Container Platform on user-provisioned infrastructure.

**IMPORTANT**

The validation steps detailed in this section must succeed before you install your cluster.

Prerequisites

- You have configured the required DNS records for your user-provisioned infrastructure.

Procedure

- From your installation node, run DNS lookups against the record names of the Kubernetes API, the wildcard routes, and the cluster nodes. Validate that the IP addresses contained in the responses correspond to the correct components.
 - Perform a lookup against the Kubernetes API record name. Check that the result points to the IP address of the API load balancer:

```
$ dig +noall +answer @<nameserver_ip> api.<cluster_name>.<base_domain> 1
```

- Replace **<nameserver_ip>** with the IP address of the nameserver, **<cluster_name>** with your cluster name, and **<base_domain>** with your base domain name.

Example output

```
api.ocp4.example.com. 604800 IN A 192.168.1.5
```

- Perform a lookup against the Kubernetes internal API record name. Check that the result points to the IP address of the API load balancer:

```
$ dig +noall +answer @<nameserver_ip> api-int.<cluster_name>.<base_domain>
```

Example output

```
api-int.ocp4.example.com. 604800 IN A 192.168.1.5
```

- Test an example ***.apps.<cluster_name>.<base_domain>** DNS wildcard lookup. All of the application wildcard lookups must resolve to the IP address of the application ingress load balancer:

```
$ dig +noall +answer @<nameserver_ip> random.apps.<cluster_name>.<base_domain>
```

Example output

```
random.apps.ocp4.example.com. 604800 IN A 192.168.1.5
```



NOTE

In the example outputs, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

You can replace **random** with another wildcard value. For example, you can query the route to the OpenShift Container Platform console:

```
$ dig +noall +answer @<nameserver_ip> console-openshift-console.apps.<cluster_name>.<base_domain>
```

Example output

```
console-openshift-console.apps.ocp4.example.com. 604800 IN A 192.168.1.5
```

- d. Run a lookup against the bootstrap DNS record name. Check that the result points to the IP address of the bootstrap node:

```
$ dig +noall +answer @<nameserver_ip> bootstrap.<cluster_name>.<base_domain>
```

Example output

```
bootstrap.ocp4.example.com. 604800 IN A 192.168.1.96
```

- e. Use this method to perform lookups against the DNS record names for the control plane and compute nodes. Check that the results correspond to the IP addresses of each node.
2. From your installation node, run reverse DNS lookups against the IP addresses of the load balancer and the cluster nodes. Validate that the record names contained in the responses correspond to the correct components.
 - a. Perform a reverse lookup against the IP address of the API load balancer. Check that the response includes the record names for the Kubernetes API and the Kubernetes internal API:

```
$ dig +noall +answer @<nameserver_ip> -x 192.168.1.5
```

Example output

```
5.1.168.192.in-addr.arpa. 604800 IN PTR api-int.ocp4.example.com. 1  
5.1.168.192.in-addr.arpa. 604800 IN PTR api.ocp4.example.com. 2
```

1 Provides the record name for the Kubernetes internal API.

2 Provides the record name for the Kubernetes API.

**NOTE**

A PTR record is not required for the OpenShift Container Platform application wildcard. No validation step is needed for reverse DNS resolution against the IP address of the application ingress load balancer.

- b. Perform a reverse lookup against the IP address of the bootstrap node. Check that the result points to the DNS record name of the bootstrap node:

```
$ dig +noall +answer @<nameserver_ip> -x 192.168.1.96
```

Example output

```
96.1.168.192.in-addr.arpa. 604800 IN PTR bootstrap.ocp4.example.com.
```

- c. Use this method to perform reverse lookups against the IP addresses for the control plane and compute nodes. Check that the results correspond to the DNS record names of each node.

18.6.6. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The **./openshift-install gather** command also requires the SSH public key to be in place on the cluster nodes.

**IMPORTANT**


Do not skip this procedure in production environments, where disaster recovery and debugging is required.

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1** Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.

**NOTE**

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.


2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.

**NOTE**

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```

**NOTE**

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

18.6.7. Obtaining the installation program


Before you install OpenShift Container Platform, download the installation file on the host you are using for installation.

Prerequisites

- You have a computer that runs Linux or macOS, with 500 MB of local disk space.

Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select your infrastructure provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.



IMPORTANT

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.



IMPORTANT

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.

4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#) . This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

18.6.8. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.



IMPORTANT

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**.
To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.



NOTE

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your PATH.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

18.6.9. Manually creating the installation configuration file

Installing the cluster requires that you manually create the installation configuration file.

Prerequisites

- You have an SSH public key on your local machine to provide to the installation program. The key will be used for SSH authentication onto your cluster nodes for debugging and disaster recovery.
- You have obtained the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create an installation directory to store your required installation assets in:

```
$ mkdir <installation_directory>
```

**IMPORTANT**

You must create a directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

2. Customize the sample **install-config.yaml** file template that is provided and save it in the **<installation_directory>**.

**NOTE**

You must name this configuration file **install-config.yaml**.

3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.

**IMPORTANT**

The **install-config.yaml** file is consumed during the next step of the installation process. You must back it up now.

Additional resources

- [Installation configuration parameters for IBM Z®](#)

18.6.9.1. Sample install-config.yaml file for IBM Z

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.

```

apiVersion: v1
baseDomain: example.com 1
compute: 2
- hyperthreading: Enabled 3
  name: worker
  replicas: 0 4
  architecture: s390x
controlPlane: 5
  hyperthreading: Enabled 6
  name: master
  replicas: 3 7
  architecture: s390x
metadata:
  name: test 8
networking:
  clusterNetwork:
  - cidr: 10.128.0.0/14 9
    hostPrefix: 23 10
  networkType: OVNKubernetes 11
  serviceNetwork: 12


```

```

- 172.30.0.0/16
platform:
  none: {} 13
  fips: false 14
  pullSecret: '{"auths": ...}' 15
  sshKey: 'ssh-ed25519 AAAA...' 16

```

- 1** The base domain of the cluster. All DNS records must be sub-domains of this base and include the cluster name.
- 2** **5** The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.
- 3** **6** Specifies whether to enable or disable simultaneous multithreading (SMT), or hyperthreading. By default, SMT is enabled to increase the performance of the cores in your machines. You can disable it by setting the parameter value to **Disabled**. If you disable SMT, you must disable it in all cluster machines; this includes both control plane and compute machines.



NOTE

Simultaneous multithreading (SMT) is enabled by default. If SMT is not available on your OpenShift Container Platform nodes, the **hyperthreading** parameter has no effect.



IMPORTANT

If you disable **hyperthreading**, whether on your OpenShift Container Platform nodes or in the **install-config.yaml** file, ensure that your capacity planning accounts for the dramatically decreased machine performance.

- 4** You must set this value to **0** when you install OpenShift Container Platform on user-provisioned infrastructure. In installer-provisioned installations, the parameter controls the number of compute machines that the cluster creates and manages for you. In user-provisioned installations, you must manually deploy the compute machines before you finish installing the cluster.



NOTE

If you are installing a three-node cluster, do not deploy any compute machines when you install the Red Hat Enterprise Linux CoreOS (RHCOS) machines.

- 7** The number of control plane machines that you add to the cluster. Because the cluster uses these values as the number of etcd endpoints in the cluster, the value must match the number of control plane machines that you deploy.
- 8** The cluster name that you specified in your DNS records.
- 9** A block of IP addresses from which pod IP addresses are allocated. This block must not overlap with existing physical networks. These IP addresses are used for the pod network. If you need to access the pods from an external network, you must configure load balancers and routers to manage the traffic.

**NOTE**

Class E CIDR range is reserved for a future use. To use the Class E CIDR range, you must ensure your networking environment accepts the IP addresses within the Class E CIDR range.

- 10 The subnet prefix length to assign to each individual node. For example, if **hostPrefix** is set to **23**, then each node is assigned a **/23** subnet out of the given **cidr**, which allows for 510 ($2^{(32 - 23)} - 2$) pod IP addresses. If you are required to provide access to nodes from an external network, configure load balancers and routers to manage the traffic.
- 11 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 12 The IP address pool to use for service IP addresses. You can enter only one IP address pool. This block must not overlap with existing physical networks. If you need to access the services from an external network, configure load balancers and routers to manage the traffic.
- 13 You must set the platform to **none**. You cannot provide additional platform configuration variables for IBM Z® infrastructure.

**IMPORTANT**

Clusters that are installed with the platform type **none** are unable to use some features, such as managing compute machines with the Machine API. This limitation applies even if the compute machines that are attached to the cluster are installed on a platform that would normally support the feature. This parameter cannot be changed after installation.


- 14 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.

**IMPORTANT**

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 15 The [pull secret from Red Hat OpenShift Cluster Manager](#) . This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.
- 16 The SSH public key for the **core** user in Red Hat Enterprise Linux CoreOS (RHCOS).

**NOTE**

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

18.6.9.2. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.

**NOTE**

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
additionalTrustBundle: | 4
  -----BEGIN CERTIFICATE-----
  <MY_TRUSTED_CA_CERT>
  -----END CERTIFICATE-----
additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.

2 A proxy URL to use for creating HTTPS connections outside the cluster.

- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with `.` to match subdomains only. For
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.

**NOTE**

The installation program does not support the proxy **readinessEndpoints** field.

**NOTE**

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.

**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

18.6.9.3. Configuring a three-node cluster

Optionally, you can deploy zero compute machines in a minimal three node cluster that consists of three control plane machines only. This provides smaller, more resource efficient clusters for cluster administrators and developers to use for testing, development, and production.

In three-node OpenShift Container Platform environments, the three control plane machines are schedulable, which means that your application workloads are scheduled to run on them.

Prerequisites

- You have an existing **install-config.yaml** file.

Procedure

- Ensure that the number of compute replicas is set to **0** in your **install-config.yaml** file, as shown in the following **compute** stanza:

```
compute:
- name: worker
  platform: {}
  replicas: 0
```



NOTE

You must set the value of the **replicas** parameter for the compute machines to **0** when you install OpenShift Container Platform on user-provisioned infrastructure, regardless of the number of compute machines you are deploying. In installer-provisioned installations, the parameter controls the number of compute machines that the cluster creates and manages for you. This does not apply to user-provisioned installations, where the compute machines are deployed manually.



NOTE

The preferred resource for control plane nodes is six vCPUs and 21 GB. For three control plane nodes this is the memory + vCPU equivalent of a minimum five-node cluster. You should back the three nodes, each installed on a 120 GB disk, with three IFLs that are SMT2 enabled. The minimum tested setup is three vCPUs and 10 GB on a 120 GB disk for each control plane node.

For three-node cluster installations, follow these next steps:

- If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes. See the *Load balancing requirements for user-provisioned infrastructure* section for more information.
- When you create the Kubernetes manifest files in the following procedure, ensure that the **mastersSchedulable** parameter in the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** file is set to **true**. This enables your application workloads to run on the control plane nodes.
- Do not deploy any compute nodes when you create the Red Hat Enterprise Linux CoreOS (RHCOS) machines.

18.6.10. Cluster Network Operator configuration

The configuration for the cluster network is specified as part of the Cluster Network Operator (CNO) configuration and stored in a custom resource (CR) object that is named **cluster**. The CR specifies the fields for the **Network** API in the **operator.openshift.io** API group.

The CNO configuration inherits the following fields during cluster installation from the **Network** API in the **Network.config.openshift.io** API group:

clusterNetwork

IP address pools from which pod IP addresses are allocated.

serviceNetwork

IP address pool for services.

defaultNetwork.type

Cluster network plugin. **OVNKubernetes** is the only supported plugin during installation.

You can specify the cluster network plugin configuration for your cluster by setting the fields for the **defaultNetwork** object in the CNO object named **cluster**.

18.6.10.1. Cluster Network Operator configuration object

The fields for the Cluster Network Operator (CNO) are described in the following table:


Table 18.84. Cluster Network Operator configuration object

Field	Type	Description
metadata.name	string	The name of the CNO object. This name is always cluster .
spec.clusterNetwork	array	A list specifying the blocks of IP addresses from which pod IP addresses are allocated and the subnet prefix length assigned to each individual node in the cluster. For example: <pre>spec: clusterNetwork: - cidr: 10.128.0.0/19 hostPrefix: 23 - cidr: 10.128.32.0/19 hostPrefix: 23</pre>
spec.serviceNetwork	array	A block of IP addresses for services. The OpenShift SDN and OVN-Kubernetes network plugins support only a single IP address block for the service network. For example: <pre>spec: serviceNetwork: - 172.30.0.0/14</pre> <p>You can customize this field only in the install-config.yaml file before you create the manifests. The value is read-only in the manifest file.</p>
spec.defaultNetwork	object	Configures the network plugin for the cluster network.
spec.kubeProxyConfig	object	The fields for this object specify the kube-proxy configuration. If you are using the OVN-Kubernetes cluster network plugin, the kube-proxy configuration has no effect.

defaultNetwork object configuration

The values for the **defaultNetwork** object are defined in the following table:

Table 18.85. defaultNetwork object

Field	Type	Description
type	string	<p>OVNKubernetes. The Red Hat OpenShift Networking network plugin is selected during installation. This value cannot be changed after cluster installation.</p>  <p>NOTE</p> <p>OpenShift Container Platform uses the OVN-Kubernetes network plugin by default. OpenShift SDN is no longer available as an installation choice for new clusters.</p>
ovnKubernetesConfig	object	This object is only valid for the OVN-Kubernetes network plugin.

Configuration for the OVN-Kubernetes network plugin

The following table describes the configuration fields for the OVN-Kubernetes network plugin:

Table 18.86. ovnKubernetesConfig object

Field	Type	Description
mtu	integer	<p>The maximum transmission unit (MTU) for the Geneve (Generic Network Virtualization Encapsulation) overlay network. This is detected automatically based on the MTU of the primary network interface. You do not normally need to override the detected MTU.</p> <p>If the auto-detected value is not what you expect it to be, confirm that the MTU on the primary network interface on your nodes is correct. You cannot use this option to change the MTU value of the primary network interface on the nodes.</p> <p>If your cluster requires different MTU values for different nodes, you must set this value to 100 less than the lowest MTU value in your cluster. For example, if some nodes in your cluster have an MTU of 9001, and some have an MTU of 1500, you must set this value to 1400.</p>
genevePort	integer	The port to use for all Geneve packets. The default value is 6081 . This value cannot be changed after cluster installation.
ipsecConfig	object	Specify a configuration object for customizing the IPsec configuration.
ipv4	object	Specifies a configuration object for IPv4 settings.


Field	Type	Description
ipv6	object	Specifies a configuration object for IPv6 settings.
policyAuditConfig	object	Specify a configuration object for customizing network policy audit logging. If unset, the defaults audit log settings are used.
gatewayConfig	object	Optional: Specify a configuration object for customizing how egress traffic is sent to the node gateway.  NOTE While migrating egress traffic, you can expect some disruption to workloads and service traffic until the Cluster Network Operator (CNO) successfully rolls out the changes.

Table 18.87. `ovnKubernetesConfig.ipv4` object

Field	Type	Description
internalTransitSwitchSubnet	string	If your existing network infrastructure overlaps with the 100.88.0.0/16 IPv4 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. The subnet for the distributed transit switch that enables east-west traffic. This subnet cannot overlap with any other subnets used by OVN-Kubernetes or on the host itself. It must be large enough to accommodate one IP address per node in your cluster. The default value is 100.88.0.0/16 .
internalJoinSubnet	string	If your existing network infrastructure overlaps with the 100.64.0.0/16 IPv4 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. You must ensure that the IP address range does not overlap with any other subnet used by your OpenShift Container Platform installation. The IP address range must be larger than the maximum number of nodes that can be added to the cluster. For example, if the clusterNetwork.cidr value is 10.128.0.0/14 and the clusterNetwork.hostPrefix value is /23 , then the maximum number of nodes is 2⁽²³⁻¹⁴⁾=512 . The default value is 100.64.0.0/16 .

Table 18.88. `ovnKubernetesConfig.ipv6` object

Field	Type	Description
-------	------	-------------

Field	Type	Description
internalTransitSwitchSubnet	string	<p>If your existing network infrastructure overlaps with the fd97::/64 IPv6 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. The subnet for the distributed transit switch that enables east-west traffic. This subnet cannot overlap with any other subnets used by OVN-Kubernetes or on the host itself. It must be large enough to accommodate one IP address per node in your cluster.</p> <p>The default value is fd97::/64.</p>
internalJoinSubnet	string	<p>If your existing network infrastructure overlaps with the fd98::/64 IPv6 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. You must ensure that the IP address range does not overlap with any other subnet used by your OpenShift Container Platform installation. The IP address range must be larger than the maximum number of nodes that can be added to the cluster.</p> <p>The default value is fd98::/64.</p>

Table 18.89. `policyAuditConfig` object

Field	Type	Description
rateLimit	integer	The maximum number of messages to generate every second per node. The default value is 20 messages per second.
maxFileSize	integer	The maximum size for the audit log in bytes. The default value is 50000000 or 50 MB.
maxLogFiles	integer	The maximum number of log files that are retained.
destination	string	<p>One of the following additional audit log targets:</p> <p>libc The libc syslog() function of the journald process on the host.</p> <p>udp:<host>:<port> A syslog server. Replace <host>:<port> with the host and port of the syslog server.</p> <p>unix:<file> A Unix Domain Socket file specified by <file>.</p> <p>null Do not send the audit logs to any additional target.</p>
syslogFacility	string	The syslog facility, such as kern , as defined by RFC5424. The default value is local0 .

Table 18.90. gatewayConfig object

Field	Type	Description
routingViaHost	boolean	<p>Set this field to true to send egress traffic from pods to the host networking stack. For highly-specialized installations and applications that rely on manually configured routes in the kernel routing table, you might want to route egress traffic to the host networking stack. By default, egress traffic is processed in OVN to exit the cluster and is not affected by specialized routes in the kernel routing table. The default value is false.</p> <p>This field has an interaction with the Open vSwitch hardware offloading feature. If you set this field to true, you do not receive the performance benefits of the offloading because egress traffic is processed by the host networking stack.</p>
ipForwarding	object	<p>You can control IP forwarding for all traffic on OVN-Kubernetes managed interfaces by using the ipForwarding specification in the Network resource. Specify Restricted to only allow IP forwarding for Kubernetes related traffic. Specify Global to allow forwarding of all IP traffic. For new installations, the default is Restricted. For updates to OpenShift Container Platform 4.14 or later, the default is Global.</p>
ipv4	object	<p>Optional: Specify an object to configure the internal OVN-Kubernetes masquerade address for host to service traffic for IPv4 addresses.</p>
ipv6	object	<p>Optional: Specify an object to configure the internal OVN-Kubernetes masquerade address for host to service traffic for IPv6 addresses.</p>

Table 18.91. gatewayConfig.ipv4 object

Field	Type	Description
internalMasqueradeSubnet	string	<p>The masquerade IPv4 addresses that are used internally to enable host to service traffic. The host is configured with these IP addresses as well as the shared gateway bridge interface. The default value is 169.254.169.0/29.</p>

Table 18.92. gatewayConfig.ipv6 object

Field	Type	Description
internalMasqueradeSubnet	string	<p>The masquerade IPv6 addresses that are used internally to enable host to service traffic. The host is configured with these IP addresses as well as the shared gateway bridge interface. The default value is fd69::/125.</p>

Table 18.93. `ipsecConfig` object


Field	Type	Description
<code>mode</code>	<code>string</code>	<p>Specifies the behavior of the IPsec implementation. Must be one of the following values:</p> <ul style="list-style-type: none"> ● Disabled: IPsec is not enabled on cluster nodes. ● External: IPsec is enabled for network traffic with external hosts. ● Full: IPsec is enabled for pod traffic and network traffic with external hosts.

Example OVN-Kubernetes configuration with IPsec enabled

```

defaultNetwork:
  type: OVNKubernetes
  ovnKubernetesConfig:
    mtu: 1400
    genevePort: 6081
    ipsecConfig:
      mode: Full

```




IMPORTANT

Using OVNKubernetes can lead to a stack exhaustion problem on IBM Power®.

kubeProxyConfig object configuration (OpenShiftSDN container network interface only)

The values for the `kubeProxyConfig` object are defined in the following table:

Table 18.94. `kubeProxyConfig` object

Field	Type	Description
<code>iptablesSyncPeriod</code>	<code>string</code>	<p>The refresh period for <code>iptables</code> rules. The default value is 30s. Valid suffixes include s, m, and h and are described in the Go time package documentation.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p>NOTE</p> <p>Because of performance improvements introduced in OpenShift Container Platform 4.3 and greater, adjusting the <code>iptablesSyncPeriod</code> parameter is no longer necessary.</p> </div> </div>

Field	Type	Description
proxyArguments.iptables-min-sync-period	array	<p>The minimum duration before refreshing iptables rules. This field ensures that the refresh does not happen too frequently. Valid suffixes include s, m, and h and are described in the Go time package. The default value is:</p> <pre>kubeProxyConfig: proxyArguments: iptables-min-sync-period: - 0s</pre>

18.6.11. Creating the Kubernetes manifest and Ignition config files

Because you must modify some cluster definition files and manually start the cluster machines, you must generate the Kubernetes manifest and Ignition config files that the cluster needs to configure the machines.

The installation configuration file transforms into the Kubernetes manifests. The manifests wrap into the Ignition configuration files, which are later used to configure the cluster machines.



IMPORTANT

- The Ignition config files that the OpenShift Container Platform installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.



NOTE

The installation program that generates the manifest and Ignition files is architecture specific and can be obtained from the [client image mirror](#). The Linux version of the installation program runs on s390x only. This installer program is also available as a Mac OS version.

Prerequisites

- You obtained the OpenShift Container Platform installation program.
- You created the **install-config.yaml** installation configuration file.

Procedure

1. Change to the directory that contains the OpenShift Container Platform installation program and generate the Kubernetes manifests for the cluster:


```
$ ./openshift-install create manifests --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the installation directory that contains the **install-config.yaml** file you created.



WARNING

If you are installing a three-node cluster, skip the following step to allow the control plane nodes to be schedulable.



IMPORTANT

When you configure control plane nodes from the default unschedulable to schedulable, additional subscriptions are required. This is because control plane nodes then become compute nodes.

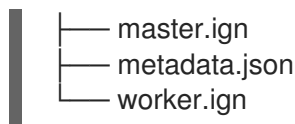
2. Check that the **mastersSchedulable** parameter in the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** Kubernetes manifest file is set to **false**. This setting prevents pods from being scheduled on the control plane machines:
 - a. Open the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** file.
 - b. Locate the **mastersSchedulable** parameter and ensure that it is set to **false**.
 - c. Save and exit the file.
3. To create the Ignition configuration files, run the following command from the directory that contains the installation program:

```
$ ./openshift-install create ignition-configs --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the same installation directory.

Ignition config files are created for the bootstrap, control plane, and compute nodes in the installation directory. The **kubeadmin-password** and **kubeconfig** files are created in the **./<installation_directory>/auth** directory:

```
.
├── auth
│   ├── kubeadmin-password
│   └── kubeconfig
└── bootstrap.ign
```

18.6.12. Configuring NBDE with static IP in an IBM Z or IBM LinuxONE environment

Enabling NBDE disk encryption in an IBM Z® or IBM® LinuxONE environment requires additional steps, which are described in detail in this section.

Prerequisites

- You have set up the External Tang Server. See [Network-bound disk encryption](#) for instructions.
- You have installed the **butane** utility.
- You have reviewed the instructions for how to create machine configs with Butane.

Procedure

1. Create Butane configuration files for the control plane and compute nodes.
The following example of a Butane configuration for a control plane node creates a file named **master-storage.bu** for disk encryption:

```

variant: openshift
version: 4.16.0
metadata:
  name: master-storage
  labels:
    machineconfiguration.openshift.io/role: master
storage:
  luks:
    - clevis:
      tang:
        - thumbprint: QcPr_NHFJammnRCA3fFMVdNBwjs
          url: http://clevis.example.com:7500
      options: ❶
        - --cipher
        - aes-cbc-essiv:sha256
    device: /dev/disk/by-partlabel/root ❷
    label: luks-root
    name: root
    wipe_volume: true
filesystems:
  - device: /dev/mapper/root
    format: xfs
    label: root
    wipe_filesystem: true
openshift:
  fips: true ❸

```

❶ The cipher option is only required if FIPS mode is enabled. Omit the entry if FIPS is disabled.

❷ For installations on DASD-type disks, replace with **device: /dev/disk/by-label/root**.

- 3 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that

2. Create a customized initramfs file to boot the machine, by running the following command:

```
$ coreos-installer pxe customize \
  /root/rhcos-bootfiles/rhcos-<release>-live-initramfs.s390x.img \
  --dest-device /dev/disk/by-id/scsi-<serial_number> --dest-karg-append \
  ip=<ip_address>::<gateway_ip>:<subnet_mask>::<network_device>:none \
  --dest-karg-append nameserver=<nameserver_ip> \
  --dest-karg-append rd.neednet=1 -o \
  /root/rhcos-bootfiles/<node_name>-initramfs.s390x.img
```



NOTE

Before first boot, you must customize the initramfs for each node in the cluster, and add PXE kernel parameters.

3. Create a parameter file that includes **ignition.platform.id=metal** and **ignition.firstboot**.

Example kernel parameter file for the control plane machine:

```
rd.neednet=1 \
console=ttySclp0 \
coreos.inst.install_dev=/dev/dasda \ 1
ignition.firstboot ignition.platform.id=metal \
coreos.live.rootfs_url=http://<http_server>/rhcos-<version>-live-rootfs.<architecture>.img \ 2
coreos.inst.ignition_url=http://<http_server>/master.ign \ 3
ip=10.19.17.2::10.19.17.1:255.255.255.0::enbddd0:none nameserver=10.19.17.1 \
zfcf.allow_lun_scan=0 \ 4
rd.znet=qeth,0.0.bddd0,0.0.bddd1,0.0.bddd2,layer2=1 \
rd.zfcf=0.0.5677,0x600606680g7f0056,0x034F000000000000 \ 5
```

- 1 For installations on DASD-type disks, add **coreos.inst.install_dev=/dev/dasda**. Omit this value for FCP-type disks.
- 2 Specify the location of the **rootfs** artifact for the **kernel** and **initramfs** you are booting. Only HTTP and HTTPS protocols are supported.
- 3 Specify the location of the Ignition config file. Use **master.ign** or **worker.ign**. Only HTTP and HTTPS protocols are supported.
- 4 For installations on FCP-type disks, add **zfcf.allow_lun_scan=0**. Omit this value for DASD-type disks.
- 5 For installations on DASD-type disks, replace with **rd.dasd=0.0.3490** to specify the DASD device.



NOTE

Write all options in the parameter file as a single line and make sure you have no newline characters.

Additional resources

- [Creating machine configs with Butane](#)

18.6.13. Installing RHCOS and starting the OpenShift Container Platform bootstrap process

To install OpenShift Container Platform on IBM Z® infrastructure that you provision, you must install Red Hat Enterprise Linux CoreOS (RHCOS) in an LPAR. When you install RHCOS, you must provide the Ignition config file that was generated by the OpenShift Container Platform installation program for the type of machine you are installing. If you have configured suitable networking, DNS, and load balancing infrastructure, the OpenShift Container Platform bootstrap process begins automatically after the RHCOS guest machines have rebooted.

Complete the following steps to create the machines.

Prerequisites

- An HTTP or HTTPS server running on your provisioning machine that is accessible to the machines you create.
- If you want to enable secure boot, you have obtained the appropriate Red Hat Product Signing Key and read [Secure boot on IBM Z and IBM LinuxONE](#) in IBM documentation.

Procedure

1. Log in to Linux on your provisioning machine.
2. Obtain the Red Hat Enterprise Linux CoreOS (RHCOS) kernel, initramfs, and rootfs files from the [RHCOS image mirror](#).



IMPORTANT

The RHCOS images might not change with every release of OpenShift Container Platform. You must download images with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Only use the appropriate kernel, initramfs, and rootfs artifacts described in the following procedure.

The file names contain the OpenShift Container Platform version number. They resemble the following examples:

- kernel: **rhcos-`<version>`-live-kernel-`<architecture>`**
- initramfs: **rhcos-`<version>`-live-initramfs.`<architecture>`.img**
- rootfs: **rhcos-`<version>`-live-rootfs.`<architecture>`.img**



NOTE

The rootfs image is the same for FCP and DASD.

3. Create parameter files. The following parameters are specific for a particular virtual machine:

- For **ip=**, specify the following seven entries:
 - i. The IP address for the machine.
 - ii. An empty string.
 - iii. The gateway.
 - iv. The netmask.
 - v. The machine host and domain name in the form **hostname.domainname**. Omit this value to let RHCOS decide.
 - vi. The network interface name. Omit this value to let RHCOS decide.
 - vii. If you use static IP addresses, specify **none**.
- For **coreos.inst.ignition_url=**, specify the Ignition file for the machine role. Use **bootstrap.ign**, **master.ign**, or **worker.ign**. Only HTTP and HTTPS protocols are supported.
- For **coreos.live.rootfs_url=**, specify the matching rootfs artifact for the kernel and initramfs you are booting. Only HTTP and HTTPS protocols are supported.
- Optional: To enable secure boot, add **coreos.inst.secure_ipl**
- For installations on DASD-type disks, complete the following tasks:
 - i. For **coreos.inst.install_dev=**, specify **/dev/dasda**.
 - ii. Use **rd.dasd=** to specify the DASD where RHCOS is to be installed.
 - iii. Leave all other parameters unchanged.

Example parameter file, **bootstrap-0.parm**, for the bootstrap machine:

```
rd.neednet=1 \
console=ttysclp0 \
coreos.inst.install_dev=/dev/dasda \
coreos.live.rootfs_url=http://<http_server>/rhcos-<version>-live-rootfs.
<architecture>.img \ 1
coreos.inst.ignition_url=http://<http_server>/bootstrap.ign \ 2
coreos.inst.secure_ipl \ 3
ip=172.18.78.2::172.18.78.1:255.255.255.0::none nameserver=172.18.78.1 \
rd.znet=qeth,0.0.bdf0,0.0.bdf1,0.0.bdf2,layer2=1,portno=0 \
zfcplib.allow_lun_scan=0 \
rd.dasd=0.0.3490
```

- 1 Specify the location of the **rootfs** artifact for the **kernel** and **initramfs** you are booting. Only HTTP and HTTPS protocols are supported.
- 2 Specify the location of the Ignition config file. Use **bootstrap.ign**, **master.ign**, or **worker.ign**. Only HTTP and HTTPS protocols are supported.
- 3 Optional: To enable secure boot, add **coreos.inst.secure_ipl**.

Write all options in the parameter file as a single line and make sure you have no newline characters.

- For installations on FCP-type disks, complete the following tasks:
 - i. Use **rd.zfcp=<adapter>,<wwpn>,<lun>** to specify the FCP disk where RHCOS is to be installed. For multipathing repeat this step for each additional path.

**NOTE**

When you install with multiple paths, you must enable multipathing directly after the installation, not at a later point in time, as this can cause problems.

- ii. Set the install device as: **coreos.inst.install_dev=/dev/disk/by-id/scsi-<serial_number>**.

**NOTE**

If additional LUNs are configured with NPIV, FCP requires **zfcp.allow_lun_scan=0**. If you must enable **zfcp.allow_lun_scan=1** because you use a CSI driver, for example, you must configure your NPIV so that each node cannot access the boot partition of another node.

- iii. Leave all other parameters unchanged.

**IMPORTANT**

Additional postinstallation steps are required to fully enable multipathing. For more information, see "Enabling multipathing with kernel arguments on RHCOS" in *Postinstallation machine configuration tasks*.

The following is an example parameter file **worker-1.parm** for a compute node with multipathing:

```
rd.neednet=1 \
console=ttysclp0 \
coreos.inst.install_dev=/dev/disk/by-id/scsi-<serial_number> \
coreos.live.rootfs_url=http://<http_server>/rhcos-<version>-live-rootfs.
<architecture>.img \
coreos.inst.ignition_url=http://<http_server>/worker.ign \
ip=172.18.78.2::172.18.78.1:255.255.255.0::none nameserver=172.18.78.1 \
rd.znet=qeth,0.0.bdf0,0.0.bdf1,0.0.bdf2,layer2=1,portno=0 \
zfcp.allow_lun_scan=0 \
rd.zfcp=0.0.1987,0x50050763070bc5e3,0x4008400B00000000 \
rd.zfcp=0.0.19C7,0x50050763070bc5e3,0x4008400B00000000 \
rd.zfcp=0.0.1987,0x50050763071bc5e3,0x4008400B00000000 \
rd.zfcp=0.0.19C7,0x50050763071bc5e3,0x4008400B00000000
```

Write all options in the parameter file as a single line and make sure you have no newline characters.

4. Transfer the initramfs, kernel, parameter files, and RHCOS images to the LPAR, for example with FTP. For details about how to transfer the files with FTP and boot, see [Installing in an LPAR](#).

5. Boot the machine
6. Repeat this procedure for the other machines in the cluster.

18.6.13.1. Advanced RHCOS installation reference

This section illustrates the networking configuration and other advanced options that allow you to modify the Red Hat Enterprise Linux CoreOS (RHCOS) manual installation process. The following tables describe the kernel arguments and command-line options you can use with the RHCOS live installer and the **coreos-installer** command.

18.6.13.1.1. Networking and bonding options for ISO installations

If you install RHCOS from an ISO image, you can add kernel arguments manually when you boot the image to configure networking for a node. If no networking arguments are specified, DHCP is activated in the initramfs when RHCOS detects that networking is required to fetch the Ignition config file.



IMPORTANT

When adding networking arguments manually, you must also add the **rd.neednet=1** kernel argument to bring the network up in the initramfs.

The following information provides examples for configuring networking and bonding on your RHCOS nodes for ISO installations. The examples describe how to use the **ip=**, **nameserver=**, and **bond=** kernel arguments.



NOTE

Ordering is important when adding the kernel arguments: **ip=**, **nameserver=**, and then **bond=**.

The networking options are passed to the **dracut** tool during system boot. For more information about the networking options supported by **dracut**, see the [dracut.cmdline manual page](#).

The following examples are the networking options for ISO installation.

Configuring DHCP or static IP addresses

To configure an IP address, either use DHCP (**ip=dhcp**) or set an individual static IP address (**ip=<host_ip>**). If setting a static IP, you must then identify the DNS server IP address (**nameserver=<dns_ip>**) on each node. The following example sets:

- The node's IP address to **10.10.10.2**
- The gateway address to **10.10.10.254**
- The netmask to **255.255.255.0**
- The hostname to **core0.example.com**
- The DNS server address to **4.4.4.41**
- The auto-configuration value to **none**. No auto-configuration is required when IP networking is configured statically.

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none
nameserver=4.4.4.41
```



NOTE

When you use DHCP to configure IP addressing for the RHCOS machines, the machines also obtain the DNS server information through DHCP. For DHCP-based deployments, you can define the DNS server address that is used by the RHCOS nodes through your DHCP server configuration.

Configuring an IP address without a static hostname

You can configure an IP address without assigning a static hostname. If a static hostname is not set by the user, it will be picked up and automatically set by a reverse DNS lookup. To configure an IP address without a static hostname refer to the following example:

- The node's IP address to **10.10.10.2**
- The gateway address to **10.10.10.254**
- The netmask to **255.255.255.0**
- The DNS server address to **4.4.4.41**
- The auto-configuration value to **none**. No auto-configuration is required when IP networking is configured statically.

```
ip=10.10.10.2::10.10.10.254:255.255.255.0::enp1s0:none
nameserver=4.4.4.41
```

Specifying multiple network interfaces

You can specify multiple network interfaces by setting multiple **ip=** entries.

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none
ip=10.10.10.3::10.10.10.254:255.255.255.0:core0.example.com:enp2s0:none
```

Configuring default gateway and route

Optional: You can configure routes to additional networks by setting an **rd.route=** value.



NOTE

When you configure one or multiple networks, one default gateway is required. If the additional network gateway is different from the primary network gateway, the default gateway must be the primary network gateway.

- Run the following command to configure the default gateway:

```
ip=::10.10.10.254:::
```

- Enter the following command to configure the route for the additional network:

```
rd.route=20.20.20.0/24:20.20.20.254:enp2s0
```

Disabling DHCP on a single interface

You can disable DHCP on a single interface, such as when there are two or more network interfaces and only one interface is being used. In the example, the **enp1s0** interface has a static networking configuration and DHCP is disabled for **enp2s0**, which is not used:

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none
ip=::::core0.example.com:enp2s0:none
```

Combining DHCP and static IP configurations

You can combine DHCP and static IP configurations on systems with multiple network interfaces, for example:

```
ip=enp1s0:dhcp
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp2s0:none
```

Configuring VLANs on individual interfaces

Optional: You can configure VLANs on individual interfaces by using the **vlan=** parameter.

- To configure a VLAN on a network interface and use a static IP address, run the following command:

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp2s0.100:none
vlan=enp2s0.100:enp2s0
```

- To configure a VLAN on a network interface and to use DHCP, run the following command:

```
ip=enp2s0.100:dhcp
vlan=enp2s0.100:enp2s0
```

Providing multiple DNS servers

You can provide multiple DNS servers by adding a **nameserver=** entry for each server, for example:

```
nameserver=1.1.1.1
nameserver=8.8.8.8
```

Bonding multiple network interfaces to a single interface

Optional: You can bond multiple network interfaces to a single interface by using the **bond=** option. Refer to the following examples:

- The syntax for configuring a bonded interface is: **bond=<name>[:<network_interfaces>] [:options]**
<name> is the bonding device name (**bond0**), **<network_interfaces>** represents a comma-separated list of physical (ethernet) interfaces (**em1,em2**), and *options* is a comma-separated list of bonding options. Enter **modinfo bonding** to see available options.
- When you create a bonded interface using **bond=**, you must specify how the IP address is assigned and other information for the bonded interface.
 - To configure the bonded interface to use DHCP, set the bond's IP address to **dhcp**. For example:

```
bond=bond0:em1,em2:mode=active-backup
ip=bond0:dhcp
```


- To configure the bonded interface to use a static IP address, enter the specific IP address you want and related information. For example:

```
bond=bond0:em1,em2:mode=active-backup,fail_over_mac=1
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:bond0:none
```

Always set the **fail_over_mac=1** option in active-backup mode, to avoid problems when shared OSA/RoCE cards are used.

Bonding multiple network interfaces to a single interface

Optional: You can configure VLANs on bonded interfaces by using the **vlan=** parameter and to use DHCP, for example:

```
ip=bond0.100:dhcp
bond=bond0:em1,em2:mode=active-backup
vlan=bond0.100:bond0
```

Use the following example to configure the bonded interface with a VLAN and to use a static IP address:

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:bond0.100:none
bond=bond0:em1,em2:mode=active-backup
vlan=bond0.100:bond0
```

Using network teaming

Optional: You can use a network teaming as an alternative to bonding by using the **team=** parameter:

- The syntax for configuring a team interface is: **team=name[:network_interfaces]**
name is the team device name (**team0**) and *network_interfaces* represents a comma-separated list of physical (ethernet) interfaces (**em1, em2**).



NOTE

Teaming is planned to be deprecated when RHCOS switches to an upcoming version of RHEL. For more information, see this [Red Hat Knowledgebase Article](#).

Use the following example to configure a network team:

```
team=team0:em1,em2
ip=team0:dhcp
```

18.6.14. Waiting for the bootstrap process to complete

The OpenShift Container Platform bootstrap process begins after the cluster nodes first boot into the persistent RHCOS environment that has been installed to disk. The configuration information provided through the Ignition config files is used to initialize the bootstrap process and install OpenShift Container Platform on the machines. You must wait for the bootstrap process to complete.

Prerequisites

- You have created the Ignition config files for your cluster.
- You have configured suitable network, DNS and load balancing infrastructure.

- You have obtained the installation program and generated the Ignition config files for your cluster.
- You installed RHCOS on your cluster machines and provided the Ignition config files that the OpenShift Container Platform installation program generated.
- Your machines have direct internet access or have an HTTP or HTTPS proxy available.

Procedure

1. Monitor the bootstrap process:

```
$ ./openshift-install --dir <installation_directory> wait-for bootstrap-complete \ 1  
--log-level=info 2
```

1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.


2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Example output

```
INFO Waiting up to 30m0s for the Kubernetes API at https://api.test.example.com:6443...  
INFO API v1.29.4 up  
INFO Waiting up to 30m0s for bootstrapping to complete...  
INFO It is now safe to remove the bootstrap resources
```

The command succeeds when the Kubernetes API server signals that it has been bootstrapped on the control plane machines.

2. After the bootstrap process is complete, remove the bootstrap machine from the load balancer.



IMPORTANT

You must remove the bootstrap machine from the load balancer at this point. You can also remove or reformat the bootstrap machine itself.

18.6.15. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

18.6.16. Approving the certificate signing requests for your machines

When you add machines to a cluster, two pending certificate signing requests (CSRs) are generated for each machine that you added. You must confirm that these CSRs are approved or, if necessary, approve them yourself. The client requests must be approved first, followed by the server requests.

Prerequisites

- You added machines to your cluster.

Procedure


1. Confirm that the cluster recognizes the machines:

```
$ oc get nodes
```

Example output

```
NAME      STATUS  ROLES  AGE  VERSION
master-0  Ready   master  63m  v1.29.4
master-1  Ready   master  63m  v1.29.4
master-2  Ready   master  64m  v1.29.4
```

The output lists all of the machines that you created.



NOTE

The preceding output might not include the compute nodes, also known as worker nodes, until some CSRs are approved.

2. Review the pending CSRs and ensure that you see the client requests with the **Pending** or **Approved** status for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

NAME	AGE	REQUESTOR	CONDITION
csr-mddf5	20m	system:node:master-01.example.com	Approved,Issued
csr-z5rln	16m	system:node:worker-21.example.com	Approved,Issued

- If the CSRs were not approved, after all of the pending CSRs for the machines you added are in **Pending** status, approve the CSRs for your cluster machines:



NOTE

Because the CSRs rotate automatically, approve your CSRs within an hour of adding the machines to the cluster. If you do not approve them within an hour, the certificates will rotate, and more than two certificates will be present for each node. You must approve all of these certificates. After the client CSR is approved, the Kubelet creates a secondary CSR for the serving certificate, which requires manual approval. Then, subsequent serving certificate renewal requests are automatically approved by the **machine-approver** if the Kubelet requests a new certificate with identical parameters.



NOTE

For clusters running on platforms that are not machine API enabled, such as bare metal and other user-provisioned infrastructure, you must implement a method of automatically approving the kubelet serving certificate requests (CSRs). If a request is not approved, then the **oc exec**, **oc rsh**, and **oc logs** commands cannot succeed, because a serving certificate is required when the API server connects to the kubelet. Any operation that contacts the Kubelet endpoint requires this certificate approval to be in place. The method must watch for new CSRs, confirm that the CSR was submitted by the **node-bootstrapper** service account in the **system:node** or **system:admin** groups, and confirm the identity of the node.

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{\n"}\n{{end}}\n{{end}}' | xargs --no-run-if-empty oc adm certificate approve
```



NOTE

Some Operators might not become available until some CSRs are approved.

- Now that your client requests are approved, you must review the server requests for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

```

NAME      AGE    REQUESTOR                                CONDITION
csr-bfd72 5m26s system:node:ip-10-0-50-126.us-east-2.compute.internal
Pending
csr-c57lv 5m26s system:node:ip-10-0-95-157.us-east-2.compute.internal
Pending
...

```

5. If the remaining CSRs are not approved, and are in the **Pending** status, approve the CSRs for your cluster machines:

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{\n"}\n{{end}}\n{{end}}' | xargs oc adm certificate approve
```

6. After all client and server CSRs have been approved, the machines have the **Ready** status. Verify this by running the following command:

```
$ oc get nodes
```

Example output

```

NAME      STATUS  ROLES  AGE  VERSION
master-0  Ready   master 73m  v1.29.4
master-1  Ready   master 73m  v1.29.4
master-2  Ready   master 74m  v1.29.4
worker-0  Ready   worker 11m  v1.29.4
worker-1  Ready   worker 11m  v1.29.4

```



NOTE

It can take a few minutes after approval of the server CSRs for the machines to transition to the **Ready** status.

Additional information

- For more information on CSRs, see [Certificate Signing Requests](#).

18.6.17. Initial Operator configuration

After the control plane initializes, you must immediately configure some Operators so that they all become available.

Prerequisites

- Your control plane has initialized.

Procedure

1. Watch the cluster components come online:

```
$ watch -n5 oc get clusteroperators
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.16.0	True	False	False	19m
baremetal	4.16.0	True	False	False	37m
cloud-credential	4.16.0	True	False	False	40m
cluster-autoscaler	4.16.0	True	False	False	37m
config-operator	4.16.0	True	False	False	38m
console	4.16.0	True	False	False	26m
csi-snapshot-controller	4.16.0	True	False	False	37m
dns	4.16.0	True	False	False	37m
etcd	4.16.0	True	False	False	36m
image-registry	4.16.0	True	False	False	31m
ingress	4.16.0	True	False	False	30m
insights	4.16.0	True	False	False	31m
kube-apiserver	4.16.0	True	False	False	26m
kube-controller-manager	4.16.0	True	False	False	36m
kube-scheduler	4.16.0	True	False	False	36m
kube-storage-version-migrator	4.16.0	True	False	False	37m
machine-api	4.16.0	True	False	False	29m
machine-approver	4.16.0	True	False	False	37m
machine-config	4.16.0	True	False	False	36m
marketplace	4.16.0	True	False	False	37m
monitoring	4.16.0	True	False	False	29m
network	4.16.0	True	False	False	38m
node-tuning	4.16.0	True	False	False	37m
openshift-apiserver	4.16.0	True	False	False	32m
openshift-controller-manager	4.16.0	True	False	False	30m
openshift-samples	4.16.0	True	False	False	32m
operator-lifecycle-manager	4.16.0	True	False	False	37m
operator-lifecycle-manager-catalog	4.16.0	True	False	False	37m
operator-lifecycle-manager-packageserver	4.16.0	True	False	False	32m
service-ca	4.16.0	True	False	False	38m
storage	4.16.0	True	False	False	37m

2. Configure the Operators that are not available.

18.6.17.1. Image registry storage configuration

The Image Registry Operator is not initially available for platforms that do not provide default storage. After installation, you must configure your registry to use storage so that the Registry Operator is made available.

Instructions are shown for configuring a persistent volume, which is required for production clusters. Where applicable, instructions are shown for configuring an empty directory as the storage location, which is available for only non-production clusters.

Additional instructions are provided for allowing the image registry to use block storage types by using the **Recreate** rollout strategy during upgrades.

18.6.17.1.1. Configuring registry storage for IBM Z

As a cluster administrator, following installation you must configure your registry to use storage.

Prerequisites

- You have access to the cluster as a user with the **cluster-admin** role.
- You have a cluster on IBM Z®.
- You have provisioned persistent storage for your cluster, such as Red Hat OpenShift Data Foundation.



IMPORTANT

OpenShift Container Platform supports **ReadWriteOnce** access for image registry storage when you have only one replica. **ReadWriteOnce** access also requires that the registry uses the **Recreate** rollout strategy. To deploy an image registry that supports high availability with two or more replicas, **ReadWriteMany** access is required.

- Must have 100Gi capacity.

Procedure

1. To configure your registry to use storage, change the **spec.storage.pvc** in the **configs.imageregistry/cluster** resource.



NOTE

When you use shared storage, review your security settings to prevent outside access.

2. Verify that you do not have a registry pod:

```
$ oc get pod -n openshift-image-registry -l docker-registry=default
```

Example output

```
No resources found in openshift-image-registry namespace
```



NOTE

If you do have a registry pod in your output, you do not need to continue with this procedure.

3. Check the registry configuration:

```
$ oc edit configs.imageregistry.operator.openshift.io
```

Example output

```
storage:
  pvc:
    claim:
```

Leave the **claim** field blank to allow the automatic creation of an **image-registry-storage** PVC.

4. Check the **clusteroperator** status:

```
$ oc get clusteroperator image-registry
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
image-registry	4.16	True	False	False	6h50m

5. Ensure that your registry is set to managed to enable building and pushing of images.

- Run:

```
$ oc edit configs.imageregistry/cluster
```

Then, change the line

```
managementState: Removed
```

to

```
managementState: Managed
```

18.6.17.1.2. Configuring storage for the image registry in non-production clusters

You must configure storage for the Image Registry Operator. For non-production clusters, you can set the image registry to an empty directory. If you do so, all images are lost if you restart the registry.

Procedure

- To set the image registry storage to an empty directory:

```
$ oc patch configs.imageregistry.operator.openshift.io cluster --type merge --patch '{"spec": {"storage":{"emptyDir":{}}}'
```


**WARNING**

Configure this option for only non-production clusters.

If you run this command before the Image Registry Operator initializes its components, the **oc patch** command fails with the following error:

```
Error from server (NotFound): configs.imageregistry.operator.openshift.io "cluster" not found
```

Wait a few minutes and run the command again.

18.6.18. Completing installation on user-provisioned infrastructure

After you complete the Operator configuration, you can finish installing the cluster on infrastructure that you provide.

Prerequisites

- Your control plane has initialized.
- You have completed the initial Operator configuration.

Procedure

1. Confirm that all the cluster components are online with the following command:

```
$ watch -n5 oc get clusteroperators
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.16.0	True	False	False	19m
baremetal	4.16.0	True	False	False	37m
cloud-credential	4.16.0	True	False	False	40m
cluster-autoscaler	4.16.0	True	False	False	37m
config-operator	4.16.0	True	False	False	38m
console	4.16.0	True	False	False	26m
csi-snapshot-controller	4.16.0	True	False	False	37m
dns	4.16.0	True	False	False	37m
etcd	4.16.0	True	False	False	36m
image-registry	4.16.0	True	False	False	31m
ingress	4.16.0	True	False	False	30m
insights	4.16.0	True	False	False	31m
kube-apiserver	4.16.0	True	False	False	26m
kube-controller-manager	4.16.0	True	False	False	36m
kube-scheduler	4.16.0	True	False	False	36m
kube-storage-version-migrator	4.16.0	True	False	False	37m
machine-api	4.16.0	True	False	False	29m

machine-approver	4.16.0	True	False	False	37m
machine-config	4.16.0	True	False	False	36m
marketplace	4.16.0	True	False	False	37m
monitoring	4.16.0	True	False	False	29m
network	4.16.0	True	False	False	38m
node-tuning	4.16.0	True	False	False	37m
openshift-apiserver	4.16.0	True	False	False	32m
openshift-controller-manager	4.16.0	True	False	False	30m
openshift-samples	4.16.0	True	False	False	32m
operator-lifecycle-manager	4.16.0	True	False	False	37m
operator-lifecycle-manager-catalog	4.16.0	True	False	False	37m
operator-lifecycle-manager-packageserver	4.16.0	True	False	False	32m
service-ca	4.16.0	True	False	False	38m
storage	4.16.0	True	False	False	37m

Alternatively, the following command notifies you when all of the clusters are available. It also retrieves and displays credentials:

```
$ ./openshift-install --dir <installation_directory> wait-for install-complete 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

Example output

```
INFO Waiting up to 30m0s for the cluster to initialize...
```

The command succeeds when the Cluster Version Operator finishes deploying the OpenShift Container Platform cluster from Kubernetes API server.

IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

2. Confirm that the Kubernetes API server is communicating with the pods.

- a. To view a list of all pods, use the following command:

```
$ oc get pods --all-namespaces
```

Example output

```

NAMESPACE          NAME          READY  STATUS
RESTARTS  AGE
openshift-apiserver-operator  openshift-apiserver-operator-85cb746d55-zqhs8  1/1
Running    1    9m
openshift-apiserver          apiserver-67b9g          1/1  Running  0
3m
openshift-apiserver          apiserver-ljcmx          1/1  Running  0
1m
openshift-apiserver          apiserver-z25h4          1/1  Running  0
2m
openshift-authentication-operator  authentication-operator-69d5d8bf84-vh2n8  1/1
Running    0    5m
...

```

- b. View the logs for a pod that is listed in the output of the previous command by using the following command:

```
$ oc logs <pod_name> -n <namespace> 1
```

- 1** Specify the pod name and namespace, as shown in the output of the previous command.

If the pod logs display, the Kubernetes API server can communicate with the cluster machines.

3. For an installation with Fibre Channel Protocol (FCP), additional steps are required to enable multipathing. Do not enable multipathing during installation. See "Enabling multipathing with kernel arguments on RHCOS" in the *Postinstallation machine configuration tasks* documentation for more information.

Verification

If you have enabled secure boot during the OpenShift Container Platform bootstrap process, the following verification steps are required:

1. Debug the node by running the following command:

```
$ oc debug node/<node_name>
chroot /host
```

2. Confirm that secure boot is enabled by running the following command:

```
$ cat /sys/firmware/ipl/secure
```

Example output

```
1 1
```

- 1** The value is **1** if secure boot is enabled and **0** if secure boot is not enabled.

3. List the re-IPL configuration by running the following command:

```
# lsreipl
```

Example output for an FCP disk

```
Re-IPL type: fcp
WWPN: 0x500507630400d1e3
LUN: 0x4001400e00000000
Device: 0.0.810e
bootprog: 0
br_lba: 0
Loadparm: ""
Bootparms: ""
clear: 0
```

Example output for a DASD disk

```
for DASD output:
Re-IPL type: ccw
Device: 0.0.525d
Loadparm: ""
clear: 0
```

4. Shut down the node by running the following command:

```
sudo shutdown -h
```

5. Initiate a boot from LPAR from the Hardware Management Console (HMC). See [Initiating a secure boot from an LPAR](#) in IBM documentation.
6. When the node is back, check the secure boot status again.

18.6.19. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service
- [How to generate SOSREPORT within OpenShift4 nodes without SSH](#) .

18.6.20. Next steps

- [Enabling multipathing with kernel arguments on RHCOS](#) .

- [Customize your cluster](#).
- If necessary, you can [opt out of remote health reporting](#).


18.7. INSTALLING A CLUSTER IN AN LPAR ON IBM Z AND IBM LINUXONE IN A RESTRICTED NETWORK

In OpenShift Container Platform version 4.16, you can install a cluster in a logical partition (LPAR) on IBM Z® or IBM® LinuxONE infrastructure that you provision in a restricted network.



NOTE

While this document refers to only IBM Z®, all information in it also applies to IBM® LinuxONE.



IMPORTANT

Additional considerations exist for non-bare metal platforms. Review the information in the [guidelines for deploying OpenShift Container Platform on non-tested platforms](#) before you install an OpenShift Container Platform cluster.

18.7.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [created a mirror registry for installation in a restricted network](#) and obtained the **imageContentSources** data for your version of OpenShift Container Platform.
- Before you begin the installation process, you must move or remove any existing installation files. This ensures that the required installation files are created and updated during the installation process.



IMPORTANT

Ensure that installation steps are done from a machine with access to the installation media.

- You provisioned [persistent storage using OpenShift Data Foundation](#) or other supported storage protocols for your cluster. To deploy a private image registry, you must set up persistent storage with **ReadWriteMany** access.
- If you use a firewall and plan to use the Telemetry service, you [configured the firewall to allow the sites](#) that your cluster requires access to.



NOTE

Be sure to also review this site list if you are configuring a proxy.

18.7.2. About installations in restricted networks

In OpenShift Container Platform 4.16, you can perform an installation that does not require an active connection to the internet to obtain software components. Restricted network installations can be completed using installer-provisioned infrastructure or user-provisioned infrastructure, depending on the cloud platform to which you are installing the cluster.

If you choose to perform a restricted network installation on a cloud platform, you still require access to its cloud APIs. Some cloud functions, like Amazon Web Service's Route 53 DNS and IAM services, require internet access. Depending on your network, you might require less internet access for an installation on bare metal hardware, Nutanix, or on VMware vSphere.

To complete a restricted network installation, you must create a registry that mirrors the contents of the OpenShift image registry and contains the installation media. You can create this registry on a mirror host, which can access both the internet and your closed network, or by using other methods that meet your restrictions.



IMPORTANT

Because of the complexity of the configuration for user-provisioned installations, consider completing a standard user-provisioned infrastructure installation before you attempt a restricted network installation using user-provisioned infrastructure. Completing this test installation might make it easier to isolate and troubleshoot any issues that might arise during your installation in a restricted network.

18.7.2.1. Additional limits

Clusters in restricted networks have the following additional limitations and restrictions:

- The **ClusterVersion** status includes an **Unable to retrieve available updates** error.
- By default, you cannot use the contents of the Developer Catalog because you cannot access the required image stream tags.

18.7.3. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to obtain the images that are necessary to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.

18.7.4. Requirements for a cluster with user-provisioned infrastructure

For a cluster that contains user-provisioned infrastructure, you must deploy all of the required machines.

This section describes the requirements for deploying OpenShift Container Platform on user-provisioned infrastructure.

18.7.4.1. Required machines for cluster installation

The smallest OpenShift Container Platform clusters require the following hosts:

Table 18.95. Minimum required hosts

Hosts	Description
One temporary bootstrap machine	The cluster requires the bootstrap machine to deploy the OpenShift Container Platform cluster on the three control plane machines. You can remove the bootstrap machine after you install the cluster.
Three control plane machines	The control plane machines run the Kubernetes and OpenShift Container Platform services that form the control plane.
At least two compute machines, which are also known as worker machines.	The workloads requested by OpenShift Container Platform users run on the compute machines.



IMPORTANT

To maintain high availability of your cluster, use separate physical hosts for these cluster machines.

The bootstrap and control plane machines must use Red Hat Enterprise Linux CoreOS (RHCOS) as the operating system. However, the compute machines can choose between Red Hat Enterprise Linux CoreOS (RHCOS), Red Hat Enterprise Linux (RHEL) 8.6 and later.

Note that RHCOS is based on Red Hat Enterprise Linux (RHEL) 9.2 and inherits all of its hardware certifications and requirements. See [Red Hat Enterprise Linux technology capabilities and limits](#) .

18.7.4.2. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 18.96. Minimum resource requirements

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)
Bootstrap	RHCOS	4	16 GB	100 GB	N/A
Control plane	RHCOS	4	16 GB	100 GB	N/A
Compute	RHCOS	2	8 GB	100 GB	N/A

1. One physical core (IFL) provides two logical cores (threads) when SMT-2 is enabled. The hypervisor can provide two or more vCPUs.



NOTE

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

18.7.4.3. Minimum IBM Z system environment

You can install OpenShift Container Platform version 4.16 on the following IBM® hardware:

- IBM® z16 (all models), IBM® z15 (all models), IBM® z14 (all models)
- IBM® LinuxONE 4 (all models), IBM® LinuxONE III (all models), IBM® LinuxONE Emperor II, IBM® LinuxONE Rockhopper II



IMPORTANT

When running OpenShift Container Platform on IBM Z® without a hypervisor use the Dynamic Partition Manager (DPM) to manage your machine.

Hardware requirements

- The equivalent of six Integrated Facilities for Linux (IFL), which are SMT2 enabled, for each cluster.
- At least one network connection to both connect to the **LoadBalancer** service and to serve data for traffic outside the cluster.



NOTE

You can use dedicated or shared IFLs to assign sufficient compute resources. Resource sharing is one of the key strengths of IBM Z®. However, you must adjust capacity correctly on each hypervisor layer and ensure sufficient resources for every OpenShift Container Platform cluster.



IMPORTANT

Since the overall performance of the cluster can be impacted, the LPARs that are used to set up the OpenShift Container Platform clusters must provide sufficient compute capacity. In this context, LPAR weight management, entitlements, and CPU shares on the hypervisor level play an important role.

Operating system requirements

- Five logical partitions (LPARs)
 - Three LPARs for OpenShift Container Platform control plane machines
 - Two LPARs for OpenShift Container Platform compute machines
- One machine for the temporary OpenShift Container Platform bootstrap machine

IBM Z network connectivity requirements

To install on IBM Z® in an LPAR, you need:

- A direct-attached OSA or RoCE network adapter
- For a preferred setup, use OSA link aggregation.

Disk storage

- FICON attached disk storage (DASDs). These can be dedicated DASDs that must be formatted as CDL, which is the default. To reach the minimum required DASD size for Red Hat Enterprise Linux CoreOS (RHCOS) installations, you need extended address volumes (EAV). If available, use HyperPAV to ensure optimal performance.
- FCP attached disk storage

Storage / Main Memory

- 16 GB for OpenShift Container Platform control plane machines
- 8 GB for OpenShift Container Platform compute machines
- 16 GB for the temporary OpenShift Container Platform bootstrap machine

Additional resources

- [Processors Resource/Systems Manager Planning Guide](#) in IBM® Documentation for PR/SM mode considerations.
- [IBM Dynamic Partition Manager \(DPM\) Guide](#) in IBM® Documentation for DPM mode considerations.
- [Topics in LPAR performance](#) for LPAR weight management and entitlements.
- [Recommended host practices for IBM Z® & IBM® LinuxONE environments](#)

18.7.4.4. Preferred IBM Z system environment

Hardware requirements

- Three LPARS that each have the equivalent of six IFLs, which are SMT2 enabled, for each

cluster.

- Two network connections to both connect to the **LoadBalancer** service and to serve data for traffic outside the cluster.
- HiperSockets that are attached to a node directly as a device. To directly connect HiperSockets to a node, you must set up a gateway to the external network via a RHEL 8 guest to bridge to the HiperSockets network.

Operating system requirements

- Three LPARs for OpenShift Container Platform control plane machines.
- At least six LPARs for OpenShift Container Platform compute machines.
- One machine or LPAR for the temporary OpenShift Container Platform bootstrap machine.

IBM Z network connectivity requirements

To install on IBM Z® in an LPAR, you need:

- A direct-attached OSA or RoCE network adapter
- For a preferred setup, use OSA link aggregation.

Disk storage

- FICON attached disk storage (DASDs). These can be dedicated DASDs that must be formatted as CDL, which is the default. To reach the minimum required DASD size for Red Hat Enterprise Linux CoreOS (RHCOS) installations, you need extended address volumes (EAV). If available, use HyperPAV to ensure optimal performance.
- FCP attached disk storage

Storage / Main Memory

- 16 GB for OpenShift Container Platform control plane machines
- 8 GB for OpenShift Container Platform compute machines
- 16 GB for the temporary OpenShift Container Platform bootstrap machine

18.7.4.5. Certificate signing requests management

Because your cluster has limited access to automatic machine management when you use infrastructure that you provision, you must provide a mechanism for approving cluster certificate signing requests (CSRs) after installation. The **kube-controller-manager** only approves the kubelet client CSRs. The **machine-approver** cannot guarantee the validity of a serving certificate that is requested by using kubelet credentials because it cannot confirm that the correct machine issued the request. You must determine and implement a method of verifying the validity of the kubelet serving certificate requests and approving them.

18.7.4.6. Networking requirements for user-provisioned infrastructure

All the Red Hat Enterprise Linux CoreOS (RHCOS) machines require networking to be configured in **inittamfs** during boot to fetch their Ignition config files.

During the initial boot, the machines require an IP address configuration that is set either through a

DHCP server or statically by providing the required boot options. After a network connection is established, the machines download their Ignition config files from an HTTP or HTTPS server. The Ignition config files are then used to set the exact state of each machine. The Machine Config Operator completes more changes to the machines, such as the application of new certificates or keys, after installation.

It is recommended to use a DHCP server for long-term management of the cluster machines. Ensure that the DHCP server is configured to provide persistent IP addresses, DNS server information, and hostnames to the cluster machines.



NOTE

If a DHCP service is not available for your user-provisioned infrastructure, you can instead provide the IP networking configuration and the address of the DNS server to the nodes at RHCOS install time. These can be passed as boot arguments if you are installing from an ISO image. See the *Installing RHCOS and starting the OpenShift Container Platform bootstrap process* section for more information about static IP provisioning and advanced networking options.

The Kubernetes API server must be able to resolve the node names of the cluster machines. If the API servers and worker nodes are in different zones, you can configure a default DNS search zone to allow the API server to resolve the node names. Another supported approach is to always refer to hosts by their fully-qualified domain names in both the node objects and all DNS requests.

18.7.4.6.1. Setting the cluster node hostnames through DHCP

On Red Hat Enterprise Linux CoreOS (RHCOS) machines, the hostname is set through NetworkManager. By default, the machines obtain their hostname through DHCP. If the hostname is not provided by DHCP, set statically through kernel arguments, or another method, it is obtained through a reverse DNS lookup. Reverse DNS lookup occurs after the network has been initialized on a node and can take time to resolve. Other system services can start prior to this and detect the hostname as **localhost** or similar. You can avoid this by using DHCP to provide the hostname for each cluster node.

Additionally, setting the hostnames through DHCP can bypass any manual DNS record name configuration errors in environments that have a DNS split-horizon implementation.

18.7.4.6.2. Network connectivity requirements

You must configure the network connectivity between machines to allow OpenShift Container Platform cluster components to communicate. Each machine must be able to resolve the hostnames of all other machines in the cluster.

This section provides details about the ports that are required.

Table 18.97. Ports used for all-machine to all-machine communications

Protocol	Port	Description
ICMP	N/A	Network reachability tests
TCP	1936	Metrics

Protocol	Port	Description
	9000-9999	Host level services, including the node exporter on ports 9100-9101 and the Cluster Version Operator on port 9099 .
	10250-10259	The default ports that Kubernetes reserves
UDP	4789	VXLAN
	6081	Geneve
	9000-9999	Host level services, including the node exporter on ports 9100-9101 .
	500	IPsec IKE packets
	4500	IPsec NAT-T packets
	123	Network Time Protocol (NTP) on UDP port 123 If an external NTP time server is configured, you must open UDP port 123 .
TCP/UDP	30000-32767	Kubernetes node port
ESP	N/A	IPsec Encapsulating Security Payload (ESP)

Table 18.98. Ports used for all-machine to control plane communications

Protocol	Port	Description
TCP	6443	Kubernetes API

Table 18.99. Ports used for control plane machine to control plane machine communications

Protocol	Port	Description
TCP	2379-2380	etcd server and peer ports

NTP configuration for user-provisioned infrastructure

OpenShift Container Platform clusters are configured to use a public Network Time Protocol (NTP) server by default. If you want to use a local enterprise NTP server, or if your cluster is being deployed in a disconnected network, you can configure the cluster to use a specific time server. For more information, see the documentation for *Configuring chrony time service*.

Additional resources

- [Configuring chrony time service](#)

18.7.4.7. User-provisioned DNS requirements

In OpenShift Container Platform deployments, DNS name resolution is required for the following components:


- The Kubernetes API
- The OpenShift Container Platform application wildcard
- The bootstrap, control plane, and compute machines

Reverse DNS resolution is also required for the Kubernetes API, the bootstrap machine, the control plane machines, and the compute machines.

DNS A/AAAA or CNAME records are used for name resolution and PTR records are used for reverse name resolution. The reverse records are important because Red Hat Enterprise Linux CoreOS (RHCOS) uses the reverse records to set the hostnames for all the nodes, unless the hostnames are provided by DHCP. Additionally, the reverse records are used to generate the certificate signing requests (CSR) that OpenShift Container Platform needs to operate.

The following DNS records are required for a user-provisioned OpenShift Container Platform cluster and they must be in place before installation. In each record, **<cluster_name>** is the cluster name and **<base_domain>** is the base domain that you specify in the **install-config.yaml** file. A complete DNS record takes the form: **<component>.<cluster_name>.<base_domain>.**

Table 18.100. Required DNS records

Component	Record	Description
Kubernetes API	api.<cluster_name>.<base_domain>.	A DNS A/AAAA or CNAME record, and a DNS PTR record, to identify the API load balancer. These records must be resolvable by both clients external to the cluster and from all the nodes within the cluster.
	api-int.<cluster_name>.<base_domain>.	A DNS A/AAAA or CNAME record, and a DNS PTR record, to internally identify the API load balancer. These records must be resolvable from all the nodes within the cluster.
		 <p>IMPORTANT</p> <p>The API server must be able to resolve the worker nodes by the hostnames that are recorded in Kubernetes. If the API server cannot resolve the node names, then proxied API calls can fail, and you cannot retrieve logs from pods.</p>

Component	Record	Description
Routes	*.apps.<cluster_name>.<base_domain>.	<p>A wildcard DNS A/AAAA or CNAME record that refers to the application ingress load balancer. The application ingress load balancer targets the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default. These records must be resolvable by both clients external to the cluster and from all the nodes within the cluster.</p> <p>For example, console-openshift-console.apps.<cluster_name>.<base_domain> is used as a wildcard route to the OpenShift Container Platform console.</p>
Bootstrap machine	bootstrap.<cluster_name>.<base_domain>.	A DNS A/AAAA or CNAME record, and a DNS PTR record, to identify the bootstrap machine. These records must be resolvable by the nodes within the cluster.
Control plane machines	<control_plane><n>.<cluster_name>.<base_domain>.	DNS A/AAAA or CNAME records and DNS PTR records to identify each machine for the control plane nodes. These records must be resolvable by the nodes within the cluster.
Compute machines	<compute><n>.<cluster_name>.<base_domain>.	DNS A/AAAA or CNAME records and DNS PTR records to identify each machine for the worker nodes. These records must be resolvable by the nodes within the cluster.



NOTE

In OpenShift Container Platform 4.4 and later, you do not need to specify etcd host and SRV records in your DNS configuration.

TIP

You can use the **dig** command to verify name and reverse name resolution. See the section on *Validating DNS resolution for user-provisioned infrastructure* for detailed validation steps.

18.7.4.7.1. Example DNS configuration for user-provisioned clusters

This section provides A and PTR record configuration samples that meet the DNS requirements for deploying OpenShift Container Platform on user-provisioned infrastructure. The samples are not meant to provide advice for choosing one DNS solution over another.

In the examples, the cluster name is **ocp4** and the base domain is **example.com**.

Example DNS A record configuration for a user-provisioned cluster

The following example is a BIND zone file that shows sample A records for name resolution in a user-provisioned cluster.

Example 18.16. Sample DNS zone database

```

$TTL 1W
@ IN SOA ns1.example.com. root (
    2019070700 ; serial
    3H ; refresh (3 hours)
    30M ; retry (30 minutes)
    2W ; expiry (2 weeks)
    1W ) ; minimum (1 week)
IN NS ns1.example.com.
IN MX 10 smtp.example.com.
;
;
ns1.example.com. IN A 192.168.1.5
smtp.example.com. IN A 192.168.1.5
;
helper.example.com. IN A 192.168.1.5
helper.ocp4.example.com. IN A 192.168.1.5
;
api.ocp4.example.com. IN A 192.168.1.5 1
api-int.ocp4.example.com. IN A 192.168.1.5 2
;
*.apps.ocp4.example.com. IN A 192.168.1.5 3
;
bootstrap.ocp4.example.com. IN A 192.168.1.96 4
;
control-plane0.ocp4.example.com. IN A 192.168.1.97 5
control-plane1.ocp4.example.com. IN A 192.168.1.98 6
control-plane2.ocp4.example.com. IN A 192.168.1.99 7
;
compute0.ocp4.example.com. IN A 192.168.1.11 8
compute1.ocp4.example.com. IN A 192.168.1.7 9
;
;EOF

```

- 1 Provides name resolution for the Kubernetes API. The record refers to the IP address of the API load balancer.
- 2 Provides name resolution for the Kubernetes API. The record refers to the IP address of the API load balancer and is used for internal cluster communications.
- 3 Provides name resolution for the wildcard routes. The record refers to the IP address of the application ingress load balancer. The application ingress load balancer targets the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.

**NOTE**

In the example, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

- 4 Provides name resolution for the bootstrap machine.
- 5 6 7 Provides name resolution for the control plane machines.
- 8 9 Provides name resolution for the compute machines.

Example DNS PTR record configuration for a user-provisioned cluster

The following example BIND zone file shows sample PTR records for reverse name resolution in a user-provisioned cluster.

Example 18.17. Sample DNS zone database for reverse records

```
$TTL 1W
@ IN SOA ns1.example.com. root (
    2019070700 ; serial
    3H ; refresh (3 hours)
    30M ; retry (30 minutes)
    2W ; expiry (2 weeks)
    1W ) ; minimum (1 week)
IN NS ns1.example.com.
;
5.1.168.192.in-addr.arpa. IN PTR api.ocp4.example.com. 1
5.1.168.192.in-addr.arpa. IN PTR api-int.ocp4.example.com. 2
;
96.1.168.192.in-addr.arpa. IN PTR bootstrap.ocp4.example.com. 3
;
97.1.168.192.in-addr.arpa. IN PTR control-plane0.ocp4.example.com. 4
98.1.168.192.in-addr.arpa. IN PTR control-plane1.ocp4.example.com. 5
99.1.168.192.in-addr.arpa. IN PTR control-plane2.ocp4.example.com. 6
;
11.1.168.192.in-addr.arpa. IN PTR compute0.ocp4.example.com. 7
7.1.168.192.in-addr.arpa. IN PTR compute1.ocp4.example.com. 8
;
;EOF
```

- 1 Provides reverse DNS resolution for the Kubernetes API. The PTR record refers to the record name of the API load balancer.
- 2 Provides reverse DNS resolution for the Kubernetes API. The PTR record refers to the record name of the API load balancer and is used for internal cluster communications.
- 3 Provides reverse DNS resolution for the bootstrap machine.
- 4 5 6 Provides reverse DNS resolution for the control plane machines.
- 7 8 Provides reverse DNS resolution for the compute machines.

**NOTE**

A PTR record is not required for the OpenShift Container Platform application wildcard.

18.7.4.8. Load balancing requirements for user-provisioned infrastructure

Before you install OpenShift Container Platform, you must provision the API and application Ingress load balancing infrastructure. In production scenarios, you can deploy the API and application Ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

**NOTE**

If you want to deploy the API and application Ingress load balancers with a Red Hat Enterprise Linux (RHEL) instance, you must purchase the RHEL subscription separately.

The load balancing infrastructure must meet the following requirements:

1. **API load balancer.** Provides a common endpoint for users, both human and machine, to interact with and configure the platform. Configure the following conditions:
 - Layer 4 load balancing only. This can be referred to as Raw TCP or SSL Passthrough mode.
 - A stateless load balancing algorithm. The options vary based on the load balancer implementation.


**IMPORTANT**

Do not configure session persistence for an API load balancer. Configuring session persistence for a Kubernetes API server might cause performance issues from excess application traffic for your OpenShift Container Platform cluster and the Kubernetes API that runs inside the cluster.

Configure the following ports on both the front and back of the load balancers:

Table 18.101. API load balancer

Port	Back-end machines (pool members)	Internal	External	Description
6443	Bootstrap and control plane. You remove the bootstrap machine from the load balancer after the bootstrap machine initializes the cluster control plane. You must configure the /readyz endpoint for the API server health check probe.	X	X	Kubernetes API server
22623	Bootstrap and control plane. You remove the bootstrap machine from the load balancer after the bootstrap machine initializes the cluster control plane.	X		Machine config server

**NOTE**

The load balancer must be configured to take a maximum of 30 seconds from the time the API server turns off the **/readyz** endpoint to the removal of the API server instance from the pool. Within the time frame after **/readyz** returns an error or becomes healthy, the endpoint must have been removed or added. Probing every 5 or 10 seconds, with two successful requests to become healthy and three to become unhealthy, are well-tested values.

2. **Application Ingress load balancer.** Provides an ingress point for application traffic flowing in from outside the cluster. A working configuration for the Ingress router is required for an OpenShift Container Platform cluster.

Configure the following conditions:

- Layer 4 load balancing only. This can be referred to as Raw TCP or SSL Passthrough mode.
- A connection-based or session-based persistence is recommended, based on the options available and types of applications that will be hosted on the platform.

TIP

If the true IP address of the client can be seen by the application Ingress load balancer, enabling source IP-based session persistence can improve performance for applications that use end-to-end TLS encryption.

Configure the following ports on both the front and back of the load balancers:

Table 18.102. Application Ingress load balancer

Port	Back-end machines (pool members)	Internal	External	Description
443	The machines that run the Ingress Controller pods, compute, or worker, by default.	X	X	HTTPS traffic
80	The machines that run the Ingress Controller pods, compute, or worker, by default.	X	X	HTTP traffic

**NOTE**

If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application Ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes.

18.7.4.8.1. Example load balancer configuration for user-provisioned clusters

This section provides an example API and application Ingress load balancer configuration that meets the load balancing requirements for user-provisioned clusters. The sample is an **/etc/haproxy/haproxy.cfg** configuration for an HAProxy load balancer. The example is not meant to provide advice for choosing one load balancing solution over another.

In the example, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.



NOTE

If you are using HAProxy as a load balancer and SELinux is set to **enforcing**, you must ensure that the HAProxy service can bind to the configured TCP port by running **setsebool -P haproxy_connect_any=1**.

Example 18.18. Sample API and application Ingress load balancer configuration

```

global
  log      127.0.0.1 local2
  pidfile  /var/run/haproxy.pid
  maxconn 4000
  daemon
defaults
  mode          http
  log           global
  option        dontlognull
  option http-server-close
  option        redispatch
  retries       3
  timeout http-request 10s
  timeout queue   1m
  timeout connect 10s
  timeout client  1m
  timeout server  1m
  timeout http-keep-alive 10s
  timeout check   10s
  maxconn        3000
listen api-server-6443 1
  bind *:6443
  mode tcp
  option httpchk GET /readyz HTTP/1.0
  option log-health-checks
  balance roundrobin
  server bootstrap bootstrap.ocp4.example.com:6443 verify none check check-ssl inter 10s fall 2
  rise 3 backup 2
  server master0 master0.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
  fall 2 rise 3
  server master1 master1.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
  fall 2 rise 3
  server master2 master2.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
  fall 2 rise 3
listen machine-config-server-22623 3
  bind *:22623
  mode tcp
  server bootstrap bootstrap.ocp4.example.com:22623 check inter 1s backup 4
  server master0 master0.ocp4.example.com:22623 check inter 1s
  server master1 master1.ocp4.example.com:22623 check inter 1s
  server master2 master2.ocp4.example.com:22623 check inter 1s
listen ingress-router-443 5

```

```

bind *:443
mode tcp
balance source
server compute0 compute0.ocp4.example.com:443 check inter 1s
server compute1 compute1.ocp4.example.com:443 check inter 1s
listen ingress-router-80 6
bind *:80
mode tcp
balance source
server compute0 compute0.ocp4.example.com:80 check inter 1s
server compute1 compute1.ocp4.example.com:80 check inter 1s

```

- 1** Port **6443** handles the Kubernetes API traffic and points to the control plane machines.
- 2** **4** The bootstrap entries must be in place before the OpenShift Container Platform cluster installation and they must be removed after the bootstrap process is complete.
- 3** Port **22623** handles the machine config server traffic and points to the control plane machines.
- 5** Port **443** handles the HTTPS traffic and points to the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.
- 6** Port **80** handles the HTTP traffic and points to the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.



NOTE

If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application Ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes.

TIP

If you are using HAProxy as a load balancer, you can check that the **haproxy** process is listening on ports **6443**, **22623**, **443**, and **80** by running **netstat -nltp** on the HAProxy node.

18.7.5. Preparing the user-provisioned infrastructure

Before you install OpenShift Container Platform on user-provisioned infrastructure, you must prepare the underlying infrastructure.

This section provides details about the high-level steps required to set up your cluster infrastructure in preparation for an OpenShift Container Platform installation. This includes configuring IP networking and network connectivity for your cluster nodes, preparing a web server for the Ignition files, enabling the required ports through your firewall, and setting up the required DNS and load balancing infrastructure.


After preparation, your cluster infrastructure must meet the requirements outlined in the *Requirements for a cluster with user-provisioned infrastructure* section.

Prerequisites

- You have reviewed the [OpenShift Container Platform 4.x Tested Integrations](#) page.
- You have reviewed the infrastructure requirements detailed in the *Requirements for a cluster with user-provisioned infrastructure* section.

Procedure

1. Set up static IP addresses.
2. Set up an HTTP or HTTPS server to provide Ignition files to the cluster nodes.
3. Ensure that your network infrastructure provides the required network connectivity between the cluster components. See the *Networking requirements for user-provisioned infrastructure* section for details about the requirements.
4. Configure your firewall to enable the ports required for the OpenShift Container Platform cluster components to communicate. See *Networking requirements for user-provisioned infrastructure* section for details about the ports that are required.



IMPORTANT

By default, port **1936** is accessible for an OpenShift Container Platform cluster, because each control plane node needs access to this port.

Avoid using the Ingress load balancer to expose this port, because doing so might result in the exposure of sensitive information, such as statistics and metrics, related to Ingress Controllers.


5. Setup the required DNS infrastructure for your cluster.
 - a. Configure DNS name resolution for the Kubernetes API, the application wildcard, the bootstrap machine, the control plane machines, and the compute machines.
 - b. Configure reverse DNS resolution for the Kubernetes API, the bootstrap machine, the control plane machines, and the compute machines.
See the *User-provisioned DNS requirements* section for more information about the OpenShift Container Platform DNS requirements.
6. Validate your DNS configuration.
 - a. From your installation node, run DNS lookups against the record names of the Kubernetes API, the wildcard routes, and the cluster nodes. Validate that the IP addresses in the responses correspond to the correct components.
 - b. From your installation node, run reverse DNS lookups against the IP addresses of the load balancer and the cluster nodes. Validate that the record names in the responses correspond to the correct components.
See the *Validating DNS resolution for user-provisioned infrastructure* section for detailed DNS validation steps.
7. Provision the required API and application ingress load balancing infrastructure. See the *Load balancing requirements for user-provisioned infrastructure* section for more information about the requirements.

**NOTE**

Some load balancing solutions require the DNS name resolution for the cluster nodes to be in place before the load balancing is initialized.

18.7.6. Validating DNS resolution for user-provisioned infrastructure

You can validate your DNS configuration before installing OpenShift Container Platform on user-provisioned infrastructure.

**IMPORTANT**

The validation steps detailed in this section must succeed before you install your cluster.

Prerequisites

- You have configured the required DNS records for your user-provisioned infrastructure.

Procedure

- From your installation node, run DNS lookups against the record names of the Kubernetes API, the wildcard routes, and the cluster nodes. Validate that the IP addresses contained in the responses correspond to the correct components.
 - Perform a lookup against the Kubernetes API record name. Check that the result points to the IP address of the API load balancer:

```
$ dig +noall +answer @<nameserver_ip> api.<cluster_name>.<base_domain> 1
```

- Replace **<nameserver_ip>** with the IP address of the nameserver, **<cluster_name>** with your cluster name, and **<base_domain>** with your base domain name.

Example output

```
api.ocp4.example.com. 604800 IN A 192.168.1.5
```

- Perform a lookup against the Kubernetes internal API record name. Check that the result points to the IP address of the API load balancer:

```
$ dig +noall +answer @<nameserver_ip> api-int.<cluster_name>.<base_domain>
```

Example output

```
api-int.ocp4.example.com. 604800 IN A 192.168.1.5
```

- Test an example ***.apps.<cluster_name>.<base_domain>** DNS wildcard lookup. All of the application wildcard lookups must resolve to the IP address of the application ingress load balancer:

```
$ dig +noall +answer @<nameserver_ip> random.apps.<cluster_name>.<base_domain>
```

Example output

```
random.apps.ocp4.example.com. 604800 IN A 192.168.1.5
```



NOTE

In the example outputs, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

You can replace **random** with another wildcard value. For example, you can query the route to the OpenShift Container Platform console:

```
$ dig +noall +answer @<nameserver_ip> console-openshift-console.apps.<cluster_name>.<base_domain>
```

Example output

```
console-openshift-console.apps.ocp4.example.com. 604800 IN A 192.168.1.5
```

- d. Run a lookup against the bootstrap DNS record name. Check that the result points to the IP address of the bootstrap node:

```
$ dig +noall +answer @<nameserver_ip> bootstrap.<cluster_name>.<base_domain>
```

Example output

```
bootstrap.ocp4.example.com. 604800 IN A 192.168.1.96
```

- e. Use this method to perform lookups against the DNS record names for the control plane and compute nodes. Check that the results correspond to the IP addresses of each node.
2. From your installation node, run reverse DNS lookups against the IP addresses of the load balancer and the cluster nodes. Validate that the record names contained in the responses correspond to the correct components.
 - a. Perform a reverse lookup against the IP address of the API load balancer. Check that the response includes the record names for the Kubernetes API and the Kubernetes internal API:

```
$ dig +noall +answer @<nameserver_ip> -x 192.168.1.5
```

Example output

```
5.1.168.192.in-addr.arpa. 604800 IN PTR api-int.ocp4.example.com. 1  
5.1.168.192.in-addr.arpa. 604800 IN PTR api.ocp4.example.com. 2
```

1 Provides the record name for the Kubernetes internal API.

2 Provides the record name for the Kubernetes API.

**NOTE**

A PTR record is not required for the OpenShift Container Platform application wildcard. No validation step is needed for reverse DNS resolution against the IP address of the application ingress load balancer.

- b. Perform a reverse lookup against the IP address of the bootstrap node. Check that the result points to the DNS record name of the bootstrap node:

```
$ dig +noall +answer @<nameserver_ip> -x 192.168.1.96
```

Example output

```
96.1.168.192.in-addr.arpa. 604800 IN PTR bootstrap.ocp4.example.com.
```


- c. Use this method to perform reverse lookups against the IP addresses for the control plane and compute nodes. Check that the results correspond to the DNS record names of each node.

18.7.7. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The **./openshift-install gather** command also requires the SSH public key to be in place on the cluster nodes.

**IMPORTANT**

Do not skip this procedure in production environments, where disaster recovery and debugging is required.

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1** Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.

**NOTE**

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.

**NOTE**

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```

**NOTE**

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

18.7.8. Manually creating the installation configuration file

Installing the cluster requires that you manually create the installation configuration file.

Prerequisites

- You have an SSH public key on your local machine to provide to the installation program. The key will be used for SSH authentication onto your cluster nodes for debugging and disaster recovery.
- You have obtained the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create an installation directory to store your required installation assets in:

```
$ mkdir <installation_directory>
```



IMPORTANT

You must create a directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

2. Customize the sample **install-config.yaml** file template that is provided and save it in the **<installation_directory>**.



NOTE

You must name this configuration file **install-config.yaml**.

3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the next step of the installation process. You must back it up now.

Additional resources

- [Installation configuration parameters for IBM Z®](#)

18.7.8.1. Sample install-config.yaml file for IBM Z

**NOTE**

Simultaneous multithreading (SMT) is enabled by default. If SMT is not available on your OpenShift Container Platform nodes, the **hyperthreading** parameter has no effect.

**IMPORTANT**


If you disable **hyperthreading**, whether on your OpenShift Container Platform nodes or in the **install-config.yaml** file, ensure that your capacity planning accounts for the dramatically decreased machine performance.

- 4 You must set this value to **0** when you install OpenShift Container Platform on user-provisioned infrastructure. In installer-provisioned installations, the parameter controls the number of compute machines that the cluster creates and manages for you. In user-provisioned installations, you must manually deploy the compute machines before you finish installing the cluster.

**NOTE**

If you are installing a three-node cluster, do not deploy any compute machines when you install the Red Hat Enterprise Linux CoreOS (RHCOS) machines.

- 7 The number of control plane machines that you add to the cluster. Because the cluster uses these values as the number of etcd endpoints in the cluster, the value must match the number of control plane machines that you deploy.
- 8 The cluster name that you specified in your DNS records.
- 9 A block of IP addresses from which pod IP addresses are allocated. This block must not overlap with existing physical networks. These IP addresses are used for the pod network. If you need to access the pods from an external network, you must configure load balancers and routers to manage the traffic.

**NOTE**

Class E CIDR range is reserved for a future use. To use the Class E CIDR range, you must ensure your networking environment accepts the IP addresses within the Class E CIDR range.


- 10 The subnet prefix length to assign to each individual node. For example, if **hostPrefix** is set to **23**, then each node is assigned a **/23** subnet out of the given **cidr**, which allows for 510 ($2^{(32 - 23)} - 2$) pod IP addresses. If you are required to provide access to nodes from an external network, configure load balancers and routers to manage the traffic.
- 11 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 12 The IP address pool to use for service IP addresses. You can enter only one IP address pool. This block must not overlap with existing physical networks. If you need to access the services from an external network, configure load balancers and routers to manage the traffic.
- 13 You must set the platform to **none**. You cannot provide additional platform configuration variables for IBM Z® infrastructure.



IMPORTANT

Clusters that are installed with the platform type **none** are unable to use some features, such as managing compute machines with the Machine API. This limitation applies even if the compute machines that are attached to the cluster are installed on a platform that would normally support the feature. This parameter cannot be changed after installation.

- 14 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.



IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 15 For **<local_registry>**, specify the registry domain name, and optionally the port, that your mirror registry uses to serve content. For example, **registry.example.com** or **registry.example.com:5000**. For **<credentials>**, specify the base64-encoded user name and password for your mirror registry.
- 16 The SSH public key for the **core** user in Red Hat Enterprise Linux CoreOS (RHCOS).



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- 17 Add the **additionalTrustBundle** parameter and value. The value must be the contents of the certificate file that you used for your mirror registry. The certificate file can be an existing, trusted certificate authority or the self-signed certificate that you generated for the mirror registry.
- 18 Provide the **imageContentSources** section according to the output of the command that you used to mirror the repository.



IMPORTANT

- When using the **oc adm release mirror** command, use the output from the **imageContentSources** section.
- When using **oc mirror** command, use the **repositoryDigestMirrors** section of the **ImageContentSourcePolicy** file that results from running the command.
- **ImageContentSourcePolicy** is deprecated. For more information see *Configuring image registry repository mirroring*.

18.7.8.2. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
additionalTrustBundle: | 4
  -----BEGIN CERTIFICATE-----
  <MY_TRUSTED_CA_CERT>
  -----END CERTIFICATE-----
additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.

**NOTE**

The installation program does not support the proxy **readinessEndpoints** field.

**NOTE**

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.

**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

18.7.8.3. Configuring a three-node cluster

Optionally, you can deploy zero compute machines in a bare metal cluster that consists of three control plane machines only. This provides smaller, more resource efficient clusters for cluster administrators and developers to use for testing, development, and production.

In three-node OpenShift Container Platform environments, the three control plane machines are schedulable, which means that your application workloads are scheduled to run on them.

Prerequisites

- You have an existing **install-config.yaml** file.

Procedure

- Ensure that the number of compute replicas is set to **0** in your **install-config.yaml** file, as shown in the following **compute** stanza:

```
compute:
- name: worker
  platform: {}
  replicas: 0
```



NOTE

You must set the value of the **replicas** parameter for the compute machines to **0** when you install OpenShift Container Platform on user-provisioned infrastructure, regardless of the number of compute machines you are deploying. In installer-provisioned installations, the parameter controls the number of compute machines that the cluster creates and manages for you. This does not apply to user-provisioned installations, where the compute machines are deployed manually.

For three-node cluster installations, follow these next steps:

- If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes. See the *Load balancing requirements for user-provisioned infrastructure* section for more information.
- When you create the Kubernetes manifest files in the following procedure, ensure that the **mastersSchedulable** parameter in the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** file is set to **true**. This enables your application workloads to run on the control plane nodes.
- Do not deploy any compute nodes when you create the Red Hat Enterprise Linux CoreOS (RHCOS) machines.

18.7.9. Cluster Network Operator configuration

The configuration for the cluster network is specified as part of the Cluster Network Operator (CNO) configuration and stored in a custom resource (CR) object that is named **cluster**. The CR specifies the fields for the **Network** API in the **operator.openshift.io** API group.

The CNO configuration inherits the following fields during cluster installation from the **Network** API in the **Network.config.openshift.io** API group:

clusterNetwork

IP address pools from which pod IP addresses are allocated.

serviceNetwork

IP address pool for services.

defaultNetwork.type

Cluster network plugin. **OVNKubernetes** is the only supported plugin during installation.

You can specify the cluster network plugin configuration for your cluster by setting the fields for the **defaultNetwork** object in the CNO object named **cluster**.

18.7.9.1. Cluster Network Operator configuration object

The fields for the Cluster Network Operator (CNO) are described in the following table:


Table 18.103. Cluster Network Operator configuration object

Field	Type	Description
metadata.name	string	The name of the CNO object. This name is always cluster .
spec.clusterNetwork	array	<p>A list specifying the blocks of IP addresses from which pod IP addresses are allocated and the subnet prefix length assigned to each individual node in the cluster. For example:</p> <pre>spec: clusterNetwork: - cidr: 10.128.0.0/19 hostPrefix: 23 - cidr: 10.128.32.0/19 hostPrefix: 23</pre>
spec.serviceNetwork	array	<p>A block of IP addresses for services. The OpenShift SDN and OVN-Kubernetes network plugins support only a single IP address block for the service network. For example:</p> <pre>spec: serviceNetwork: - 172.30.0.0/14</pre> <p>You can customize this field only in the install-config.yaml file before you create the manifests. The value is read-only in the manifest file.</p>
spec.defaultNetwork	object	Configures the network plugin for the cluster network.
spec.kubeProxyConfig	object	The fields for this object specify the kube-proxy configuration. If you are using the OVN-Kubernetes cluster network plugin, the kube-proxy configuration has no effect.

defaultNetwork object configuration

The values for the **defaultNetwork** object are defined in the following table:

Table 18.104. `defaultNetwork` object

Field	Type	Description
<code>type</code>	<code>string</code>	<p>OVNKubernetes. The Red Hat OpenShift Networking network plugin is selected during installation. This value cannot be changed after cluster installation.</p> <div style="display: flex; align-items: flex-start;">  <div> <p>NOTE</p> <p>OpenShift Container Platform uses the OVN-Kubernetes network plugin by default. OpenShift SDN is no longer available as an installation choice for new clusters.</p> </div> </div>
<code>ovnKubernetesConfig</code>	<code>object</code>	This object is only valid for the OVN-Kubernetes network plugin.

Configuration for the OVN-Kubernetes network plugin

The following table describes the configuration fields for the OVN-Kubernetes network plugin:

Table 18.105. `ovnKubernetesConfig` object

Field	Type	Description
<code>mtu</code>	<code>integer</code>	<p>The maximum transmission unit (MTU) for the Geneve (Generic Network Virtualization Encapsulation) overlay network. This is detected automatically based on the MTU of the primary network interface. You do not normally need to override the detected MTU.</p> <p>If the auto-detected value is not what you expect it to be, confirm that the MTU on the primary network interface on your nodes is correct. You cannot use this option to change the MTU value of the primary network interface on the nodes.</p> <p>If your cluster requires different MTU values for different nodes, you must set this value to 100 less than the lowest MTU value in your cluster. For example, if some nodes in your cluster have an MTU of 9001, and some have an MTU of 1500, you must set this value to 1400.</p>
<code>genevePort</code>	<code>integer</code>	The port to use for all Geneve packets. The default value is 6081 . This value cannot be changed after cluster installation.
<code>ipsecConfig</code>	<code>object</code>	Specify a configuration object for customizing the IPsec configuration.
<code>ipv4</code>	<code>object</code>	Specifies a configuration object for IPv4 settings.


Field	Type	Description
ipv6	object	Specifies a configuration object for IPv6 settings.
policyAuditConfig	object	Specify a configuration object for customizing network policy audit logging. If unset, the defaults audit log settings are used.
gatewayConfig	object	Optional: Specify a configuration object for customizing how egress traffic is sent to the node gateway.
		 <p>NOTE</p> <p>While migrating egress traffic, you can expect some disruption to workloads and service traffic until the Cluster Network Operator (CNO) successfully rolls out the changes.</p>

Table 18.106. `ovnKubernetesConfig.ipv4` object

Field	Type	Description
internalTransitSwitchSubnet	string	<p>If your existing network infrastructure overlaps with the 100.88.0.0/16 IPv4 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. The subnet for the distributed transit switch that enables east-west traffic. This subnet cannot overlap with any other subnets used by OVN-Kubernetes or on the host itself. It must be large enough to accommodate one IP address per node in your cluster.</p> <p>The default value is 100.88.0.0/16.</p>
internalJoinSubnet	string	<p>If your existing network infrastructure overlaps with the 100.64.0.0/16 IPv4 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. You must ensure that the IP address range does not overlap with any other subnet used by your OpenShift Container Platform installation. The IP address range must be larger than the maximum number of nodes that can be added to the cluster. For example, if the clusterNetwork.cidr value is 10.128.0.0/14 and the clusterNetwork.hostPrefix value is /23, then the maximum number of nodes is 2⁽²³⁻¹⁴⁾=512.</p> <p>The default value is 100.64.0.0/16.</p>

Table 18.107. `ovnKubernetesConfig.ipv6` object

Field	Type	Description
-------	------	-------------

Field	Type	Description
internalTransitSwitchSubnet	string	<p>If your existing network infrastructure overlaps with the fd97::/64 IPv6 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. The subnet for the distributed transit switch that enables east-west traffic. This subnet cannot overlap with any other subnets used by OVN-Kubernetes or on the host itself. It must be large enough to accommodate one IP address per node in your cluster.</p> <p>The default value is fd97::/64.</p>
internalJoinSubnet	string	<p>If your existing network infrastructure overlaps with the fd98::/64 IPv6 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. You must ensure that the IP address range does not overlap with any other subnet used by your OpenShift Container Platform installation. The IP address range must be larger than the maximum number of nodes that can be added to the cluster.</p> <p>The default value is fd98::/64.</p>

Table 18.108. **policyAuditConfig** object

Field	Type	Description
rateLimit	integer	The maximum number of messages to generate every second per node. The default value is 20 messages per second.
maxFileSize	integer	The maximum size for the audit log in bytes. The default value is 50000000 or 50 MB.
maxLogFiles	integer	The maximum number of log files that are retained.
destination	string	<p>One of the following additional audit log targets:</p> <p>libc The libc syslog() function of the journald process on the host.</p> <p>udp:<host>:<port> A syslog server. Replace <host>:<port> with the host and port of the syslog server.</p> <p>unix:<file> A Unix Domain Socket file specified by <file>.</p> <p>null Do not send the audit logs to any additional target.</p>
syslogFacility	string	The syslog facility, such as kern , as defined by RFC5424. The default value is local0 .

Table 18.109. gatewayConfig object

Field	Type	Description
routingViaHost	boolean	<p>Set this field to true to send egress traffic from pods to the host networking stack. For highly-specialized installations and applications that rely on manually configured routes in the kernel routing table, you might want to route egress traffic to the host networking stack. By default, egress traffic is processed in OVN to exit the cluster and is not affected by specialized routes in the kernel routing table. The default value is false.</p> <p>This field has an interaction with the Open vSwitch hardware offloading feature. If you set this field to true, you do not receive the performance benefits of the offloading because egress traffic is processed by the host networking stack.</p>
ipForwarding	object	<p>You can control IP forwarding for all traffic on OVN-Kubernetes managed interfaces by using the ipForwarding specification in the Network resource. Specify Restricted to only allow IP forwarding for Kubernetes related traffic. Specify Global to allow forwarding of all IP traffic. For new installations, the default is Restricted. For updates to OpenShift Container Platform 4.14 or later, the default is Global.</p>
ipv4	object	<p>Optional: Specify an object to configure the internal OVN-Kubernetes masquerade address for host to service traffic for IPv4 addresses.</p>
ipv6	object	<p>Optional: Specify an object to configure the internal OVN-Kubernetes masquerade address for host to service traffic for IPv6 addresses.</p>

Table 18.110. gatewayConfig.ipv4 object

Field	Type	Description
internalMasqueradeSubnet	string	<p>The masquerade IPv4 addresses that are used internally to enable host to service traffic. The host is configured with these IP addresses as well as the shared gateway bridge interface. The default value is 169.254.169.0/29.</p>

Table 18.111. gatewayConfig.ipv6 object

Field	Type	Description
internalMasqueradeSubnet	string	<p>The masquerade IPv6 addresses that are used internally to enable host to service traffic. The host is configured with these IP addresses as well as the shared gateway bridge interface. The default value is fd69::/125.</p>

Table 18.112. `ipsecConfig` object

Field	Type	Description
<code>mode</code>	<code>string</code>	<p>Specifies the behavior of the IPsec implementation. Must be one of the following values:</p> <ul style="list-style-type: none"> ● Disabled: IPsec is not enabled on cluster nodes. ● External: IPsec is enabled for network traffic with external hosts. ● Full: IPsec is enabled for pod traffic and network traffic with external hosts.

Example OVN-Kubernetes configuration with IPsec enabled

```

defaultNetwork:
  type: OVNKubernetes
  ovnKubernetesConfig:
    mtu: 1400
    genevePort: 6081
    ipsecConfig:
      mode: Full

```




IMPORTANT

Using OVNKubernetes can lead to a stack exhaustion problem on IBM Power®.

kubeProxyConfig object configuration (OpenShiftSDN container network interface only)

The values for the `kubeProxyConfig` object are defined in the following table:

Table 18.113. `kubeProxyConfig` object


Field	Type	Description
<code>iptablesSyncPeriod</code>	<code>string</code>	<p>The refresh period for <code>iptables</code> rules. The default value is 30s. Valid suffixes include s, m, and h and are described in the Go time package documentation.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p>NOTE</p> <p>Because of performance improvements introduced in OpenShift Container Platform 4.3 and greater, adjusting the <code>iptablesSyncPeriod</code> parameter is no longer necessary.</p> </div> </div>

Field	Type	Description
proxyArguments.iptables-min-sync-period	array	<p>The minimum duration before refreshing iptables rules. This field ensures that the refresh does not happen too frequently. Valid suffixes include s, m, and h and are described in the Go time package. The default value is:</p> <pre>kubeProxyConfig: proxyArguments: iptables-min-sync-period: - 0s</pre>

18.7.10. Creating the Kubernetes manifest and Ignition config files

Because you must modify some cluster definition files and manually start the cluster machines, you must generate the Kubernetes manifest and Ignition config files that the cluster needs to configure the machines.

The installation configuration file transforms into the Kubernetes manifests. The manifests wrap into the Ignition configuration files, which are later used to configure the cluster machines.



IMPORTANT

- The Ignition config files that the OpenShift Container Platform installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.



NOTE

The installation program that generates the manifest and Ignition files is architecture specific and can be obtained from the [client image mirror](#). The Linux version of the installation program runs on s390x only. This installer program is also available as a Mac OS version.

Prerequisites

- You obtained the OpenShift Container Platform installation program. For a restricted network installation, these files are on your mirror host.

- You created the **install-config.yaml** installation configuration file.

Procedure

- Change to the directory that contains the OpenShift Container Platform installation program and generate the Kubernetes manifests for the cluster:

```
$ ./openshift-install create manifests --dir <installation_directory> 1
```

- For **<installation_directory>**, specify the installation directory that contains the **install-config.yaml** file you created.



WARNING

If you are installing a three-node cluster, skip the following step to allow the control plane nodes to be schedulable.



IMPORTANT

When you configure control plane nodes from the default unschedulable to schedulable, additional subscriptions are required. This is because control plane nodes then become compute nodes.


- Check that the **mastersSchedulable** parameter in the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** Kubernetes manifest file is set to **false**. This setting prevents pods from being scheduled on the control plane machines:
 - Open the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** file.
 - Locate the **mastersSchedulable** parameter and ensure that it is set to **false**.
 - Save and exit the file.
- To create the Ignition configuration files, run the following command from the directory that contains the installation program:

```
$ ./openshift-install create ignition-configs --dir <installation_directory> 1
```

- For **<installation_directory>**, specify the same installation directory.

Ignition config files are created for the bootstrap, control plane, and compute nodes in the installation directory. The **kubeadmin-password** and **kubeconfig** files are created in the **./<installation_directory>/auth** directory:

```
.
├── auth
│   ├── kubeadmin-password
```

18.7.11. Configuring NBDE with static IP in an IBM Z or IBM LinuxONE environment

Enabling NBDE disk encryption in an IBM Z® or IBM® LinuxONE environment requires additional steps, which are described in detail in this section.

Prerequisites

- You have set up the External Tang Server. See [Network-bound disk encryption](#) for instructions.
- You have installed the **butane** utility.
- You have reviewed the instructions for how to create machine configs with Butane.

Procedure

1. Create Butane configuration files for the control plane and compute nodes.
The following example of a Butane configuration for a control plane node creates a file named **master-storage.bu** for disk encryption:

```

variant: openshift
version: 4.16.0
metadata:
  name: master-storage
labels:
  machineconfiguration.openshift.io/role: master
storage:
  luks:
    - clevis:
      tang:
        - thumbprint: QcPr_NHFJammnRCA3fFMVdNBwjs
          url: http://clevis.example.com:7500
      options: ❶
        - --cipher
        - aes-cbc-essiv:sha256
    device: /dev/disk/by-partlabel/root ❷
    label: luks-root
    name: root
    wipe_volume: true
filesystems:
  - device: /dev/mapper/root
    format: xfs
    label: root
    wipe_filesystem: true
openshift:
  fips: true ❸

```

- ❶ The cipher option is only required if FIPS mode is enabled. Omit the entry if FIPS is disabled.

- 2 For installations on DASD-type disks, replace with **device: /dev/disk/by-label/root**.
- 3 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.

2. Create a customized initramfs file to boot the machine, by running the following command:

```
$ coreos-installer pxe customize \
  /root/rhcos-bootfiles/rhcos-<release>-live-initramfs.s390x.img \
  --dest-device /dev/disk/by-id/scsi-<serial_number> --dest-karg-append \
  ip=<ip_address>::<gateway_ip>:<subnet_mask>::<network_device>:none \
  --dest-karg-append nameserver=<nameserver_ip> \
  --dest-karg-append rd.neednet=1 -o \
  /root/rhcos-bootfiles/<node_name>-initramfs.s390x.img
```



NOTE

Before first boot, you must customize the initramfs for each node in the cluster, and add PXE kernel parameters.

3. Create a parameter file that includes **ignition.platform.id=metal** and **ignition.firstboot**.

Example kernel parameter file for the control plane machine:

```
rd.neednet=1 \
console=ttysclp0 \
coreos.inst.install_dev=/dev/dasda 1 \
ignition.firstboot ignition.platform.id=metal \
coreos.live.rootfs_url=http://<http_server>/rhcos-<version>-live-rootfs.<architecture>.img 2 \
coreos.inst.ignition_url=http://<http_server>/master.ign 3 \
ip=10.19.17.2::10.19.17.1:255.255.255.0::enb0d0:none nameserver=10.19.17.1 \
zfcpl.allow_lun_scan=0 4 \
rd.znet=qeth,0.0.bdd0,0.0.bdd1,0.0.bdd2,layer2=1 \
rd.zfcpl=0.0.5677,0x600606680g7f0056,0x034F000000000000 5 \
```

- 1 For installations on DASD-type disks, add **coreos.inst.install_dev=/dev/dasda**. Omit this value for FCP-type disks.
- 2 Specify the location of the **rootfs** artifact for the **kernel** and **initramfs** you are booting. Only HTTP and HTTPS protocols are supported.
- 3 Specify the location of the Ignition config file. Use **master.ign** or **worker.ign**. Only HTTP and HTTPS protocols are supported.
- 4 For installations on FCP-type disks, add **zfcpl.allow_lun_scan=0**. Omit this value for DASD-type disks.
- 5 For installations on DASD-type disks, replace with **rd.dasd=0.0.3490** to specify the DASD device.

**NOTE**

Write all options in the parameter file as a single line and make sure you have no newline characters.

Additional resources

- [Creating machine configs with Butane](#)

18.7.12. Installing RHCOS and starting the OpenShift Container Platform bootstrap process

To install OpenShift Container Platform on IBM Z® infrastructure that you provision, you must install Red Hat Enterprise Linux CoreOS (RHCOS) in an LPAR. When you install RHCOS, you must provide the Ignition config file that was generated by the OpenShift Container Platform installation program for the type of machine you are installing. If you have configured suitable networking, DNS, and load balancing infrastructure, the OpenShift Container Platform bootstrap process begins automatically after the RHCOS guest machines have rebooted.

Complete the following steps to create the machines.

Prerequisites

- An HTTP or HTTPS server running on your provisioning machine that is accessible to the machines you create.
- If you want to enable secure boot, you have obtained the appropriate Red Hat Product Signing Key and read [Secure boot on IBM Z and IBM LinuxONE](#) in IBM documentation.

Procedure


1. Log in to Linux on your provisioning machine.
2. Obtain the Red Hat Enterprise Linux CoreOS (RHCOS) kernel, initramfs, and rootfs files from the [RHCOS image mirror](#).

**IMPORTANT**

The RHCOS images might not change with every release of OpenShift Container Platform. You must download images with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Only use the appropriate kernel, initramfs, and rootfs artifacts described in the following procedure.

The file names contain the OpenShift Container Platform version number. They resemble the following examples:

- kernel: **rhcOS-<version>-live-kernel-<architecture>**
- initramfs: **rhcOS-<version>-live-initramfs.<architecture>.img**
- rootfs: **rhcOS-<version>-live-rootfs.<architecture>.img**

**NOTE**

The rootfs image is the same for FCP and DASD.

3. Create parameter files. The following parameters are specific for a particular virtual machine:

- For **ip=**, specify the following seven entries:
 - i. The IP address for the machine.
 - ii. An empty string.
 - iii. The gateway.
 - iv. The netmask.
 - v. The machine host and domain name in the form **hostname.domainname**. Omit this value to let RHCOS decide.
 - vi. The network interface name. Omit this value to let RHCOS decide.
 - vii. If you use static IP addresses, specify **none**.
- For **coreos.inst.ignition_url=**, specify the Ignition file for the machine role. Use **bootstrap.ign**, **master.ign**, or **worker.ign**. Only HTTP and HTTPS protocols are supported.
- For **coreos.live.rootfs_url=**, specify the matching rootfs artifact for the kernel and initramfs you are booting. Only HTTP and HTTPS protocols are supported.
- Optional: To enable secure boot, add **coreos.inst.secure_ign**
- For installations on DASD-type disks, complete the following tasks:
 - i. For **coreos.inst.install_dev=**, specify **/dev/dasda**.
 - ii. Use **rd.dasd=** to specify the DASD where RHCOS is to be installed.
 - iii. Leave all other parameters unchanged.

Example parameter file, **bootstrap-0.parm**, for the bootstrap machine:

```
rd.neednet=1 \
console=ttysclp0 \
coreos.inst.install_dev=/dev/dasda \
coreos.live.rootfs_url=http://<http_server>/rhcos-<version>-live-rootfs.
<architecture>.img \ 1
coreos.inst.ignition_url=http://<http_server>/bootstrap.ign \ 2
coreos.inst.secure_ign \ 3
ip=172.18.78.2::172.18.78.1:255.255.255.0::none nameserver=172.18.78.1 \
rd.znet=qeth,0.0.bdf0,0.0.bdf1,0.0.bdf2,layer2=1,portno=0 \
zfcplib.allow_lun_scan=0 \
rd.dasd=0.0.3490
```

- 1 Specify the location of the **rootfs** artifact for the **kernel** and **initramfs** you are booting. Only HTTP and HTTPS protocols are supported.
- 2 Specify the location of the Ignition config file. Use **bootstrap.ign**, **master.ign**, or

- 3 Optional: To enable secure boot, add **coreos.inst.secure_ipl**.

Write all options in the parameter file as a single line and make sure you have no newline characters.

- For installations on FCP-type disks, complete the following tasks:
 - i. Use **rd.zfcp=<adapter>,<wwpn>,<lun>** to specify the FCP disk where RHCOS is to be installed. For multipathing repeat this step for each additional path.



NOTE

When you install with multiple paths, you must enable multipathing directly after the installation, not at a later point in time, as this can cause problems.

- ii. Set the install device as: **coreos.inst.install_dev=/dev/disk/by-id/scsi-<serial_number>**.



NOTE

If additional LUNs are configured with NPIV, FCP requires **zfcp.allow_lun_scan=0**. If you must enable **zfcp.allow_lun_scan=1** because you use a CSI driver, for example, you must configure your NPIV so that each node cannot access the boot partition of another node.

- iii. Leave all other parameters unchanged.



IMPORTANT

Additional postinstallation steps are required to fully enable multipathing. For more information, see "Enabling multipathing with kernel arguments on RHCOS" in *Postinstallation machine configuration tasks*.

The following is an example parameter file **worker-1.parm** for a compute node with multipathing:

```
rd.neednet=1 \
console=ttysclp0 \
coreos.inst.install_dev=/dev/disk/by-id/scsi-<serial_number> \
coreos.live.rootfs_url=http://<http_server>/rhcos-<version>-live-rootfs.
<architecture>.img \
coreos.inst.ignition_url=http://<http_server>/worker.ign \
ip=172.18.78.2::172.18.78.1:255.255.255.0::none nameserver=172.18.78.1 \
rd.znet=qeth,0.0.bdf0,0.0.bdf1,0.0.bdf2,layer2=1,portno=0 \
zfcp.allow_lun_scan=0 \
rd.zfcp=0.0.1987,0x50050763070bc5e3,0x4008400B00000000 \
rd.zfcp=0.0.19C7,0x50050763070bc5e3,0x4008400B00000000 \
rd.zfcp=0.0.1987,0x50050763071bc5e3,0x4008400B00000000 \
rd.zfcp=0.0.19C7,0x50050763071bc5e3,0x4008400B00000000
```

Write all options in the parameter file as a single line and make sure you have no newline characters.


4. Transfer the `initramfs`, kernel, parameter files, and RHCOS images to the LPAR, for example with FTP. For details about how to transfer the files with FTP and boot, see [Installing in an LPAR](#).
5. Boot the machine
6. Repeat this procedure for the other machines in the cluster.

18.7.12.1. Advanced RHCOS installation reference

This section illustrates the networking configuration and other advanced options that allow you to modify the Red Hat Enterprise Linux CoreOS (RHCOS) manual installation process. The following tables describe the kernel arguments and command-line options you can use with the RHCOS live installer and the `coreos-installer` command.

18.7.12.1.1. Networking and bonding options for ISO installations

If you install RHCOS from an ISO image, you can add kernel arguments manually when you boot the image to configure networking for a node. If no networking arguments are specified, DHCP is activated in the `initramfs` when RHCOS detects that networking is required to fetch the Ignition config file.



IMPORTANT

When adding networking arguments manually, you must also add the `rd.neednet=1` kernel argument to bring the network up in the `initramfs`.

The following information provides examples for configuring networking and bonding on your RHCOS nodes for ISO installations. The examples describe how to use the `ip=`, `nameserver=`, and `bond=` kernel arguments.



NOTE

Ordering is important when adding the kernel arguments: `ip=`, `nameserver=`, and then `bond=`.

The networking options are passed to the `dracut` tool during system boot. For more information about the networking options supported by `dracut`, see the [dracut.cmdline manual page](#).

The following examples are the networking options for ISO installation.

Configuring DHCP or static IP addresses

To configure an IP address, either use DHCP (`ip=dhcp`) or set an individual static IP address (`ip=<host_ip>`). If setting a static IP, you must then identify the DNS server IP address (`nameserver=<dns_ip>`) on each node. The following example sets:

- The node's IP address to **10.10.10.2**
- The gateway address to **10.10.10.254**
- The netmask to **255.255.255.0**
- The hostname to **core0.example.com**

- The DNS server address to **4.4.4.41**
- The auto-configuration value to **none**. No auto-configuration is required when IP networking is configured statically.

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none
nameserver=4.4.4.41
```



NOTE

When you use DHCP to configure IP addressing for the RHCOS machines, the machines also obtain the DNS server information through DHCP. For DHCP-based deployments, you can define the DNS server address that is used by the RHCOS nodes through your DHCP server configuration.

Configuring an IP address without a static hostname

You can configure an IP address without assigning a static hostname. If a static hostname is not set by the user, it will be picked up and automatically set by a reverse DNS lookup. To configure an IP address without a static hostname refer to the following example:

- The node's IP address to **10.10.10.2**
- The gateway address to **10.10.10.254**
- The netmask to **255.255.255.0**
- The DNS server address to **4.4.4.41**
- The auto-configuration value to **none**. No auto-configuration is required when IP networking is configured statically.

```
ip=10.10.10.2::10.10.10.254:255.255.255.0::enp1s0:none
nameserver=4.4.4.41
```

Specifying multiple network interfaces

You can specify multiple network interfaces by setting multiple **ip=** entries.

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none
ip=10.10.10.3::10.10.10.254:255.255.255.0:core0.example.com:enp2s0:none
```

Configuring default gateway and route

Optional: You can configure routes to additional networks by setting an **rd.route=** value.



NOTE

When you configure one or multiple networks, one default gateway is required. If the additional network gateway is different from the primary network gateway, the default gateway must be the primary network gateway.

- Run the following command to configure the default gateway:

```
ip=::10.10.10.254:::
```

- Enter the following command to configure the route for the additional network:

```
rd.route=20.20.20.0/24:20.20.20.254:enp2s0
```

Disabling DHCP on a single interface

You can disable DHCP on a single interface, such as when there are two or more network interfaces and only one interface is being used. In the example, the **enp1s0** interface has a static networking configuration and DHCP is disabled for **enp2s0**, which is not used:

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none
ip=::::core0.example.com:enp2s0:none
```

Combining DHCP and static IP configurations

You can combine DHCP and static IP configurations on systems with multiple network interfaces, for example:

```
ip=enp1s0:dhcp
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp2s0:none
```

Configuring VLANs on individual interfaces

Optional: You can configure VLANs on individual interfaces by using the **vlan=** parameter.

- To configure a VLAN on a network interface and use a static IP address, run the following command:

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp2s0.100:none
vlan=enp2s0.100:enp2s0
```

- To configure a VLAN on a network interface and to use DHCP, run the following command:

```
ip=enp2s0.100:dhcp
vlan=enp2s0.100:enp2s0
```

Providing multiple DNS servers

You can provide multiple DNS servers by adding a **nameserver=** entry for each server, for example:

```
nameserver=1.1.1.1
nameserver=8.8.8.8
```

Bonding multiple network interfaces to a single interface

Optional: You can bond multiple network interfaces to a single interface by using the **bond=** option. Refer to the following examples:

- The syntax for configuring a bonded interface is: **bond=<name>[:<network_interfaces>][:options]**
<name> is the bonding device name (**bond0**), **<network_interfaces>** represents a comma-separated list of physical (ethernet) interfaces (**em1,em2**), and *options* is a comma-separated list of bonding options. Enter **modinfo bonding** to see available options.
- When you create a bonded interface using **bond=**, you must specify how the IP address is assigned and other information for the bonded interface.
 - To configure the bonded interface to use DHCP, set the bond's IP address to **dhcp**. For example:


```
bond=bond0:em1,em2:mode=active-backup
ip=bond0:dhcp
```

- To configure the bonded interface to use a static IP address, enter the specific IP address you want and related information. For example:

```
bond=bond0:em1,em2:mode=active-backup,fail_over_mac=1
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:bond0:none
```

Always set the **fail_over_mac=1** option in active-backup mode, to avoid problems when shared OSA/RoCE cards are used.

Bonding multiple network interfaces to a single interface

Optional: You can configure VLANs on bonded interfaces by using the **vlan=** parameter and to use DHCP, for example:

```
ip=bond0.100:dhcp
bond=bond0:em1,em2:mode=active-backup
vlan=bond0.100:bond0
```

Use the following example to configure the bonded interface with a VLAN and to use a static IP address:

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:bond0.100:none
bond=bond0:em1,em2:mode=active-backup
vlan=bond0.100:bond0
```

Using network teaming

Optional: You can use a network teaming as an alternative to bonding by using the **team=** parameter:

- The syntax for configuring a team interface is: **team=name[:network_interfaces]**
name is the team device name (**team0**) and *network_interfaces* represents a comma-separated list of physical (ethernet) interfaces (**em1, em2**).



NOTE

Teaming is planned to be deprecated when RHCOS switches to an upcoming version of RHEL. For more information, see this [Red Hat Knowledgebase Article](#).

Use the following example to configure a network team:

```
team=team0:em1,em2
ip=team0:dhcp
```

18.7.13. Waiting for the bootstrap process to complete

The OpenShift Container Platform bootstrap process begins after the cluster nodes first boot into the persistent RHCOS environment that has been installed to disk. The configuration information provided through the Ignition config files is used to initialize the bootstrap process and install OpenShift Container Platform on the machines. You must wait for the bootstrap process to complete.

Prerequisites

- You have created the Ignition config files for your cluster.
- You have configured suitable network, DNS and load balancing infrastructure.
- You have obtained the installation program and generated the Ignition config files for your cluster.
- You installed RHCOS on your cluster machines and provided the Ignition config files that the OpenShift Container Platform installation program generated.

Procedure

1. Monitor the bootstrap process:

```
$ ./openshift-install --dir <installation_directory> wait-for bootstrap-complete \ 1  
--log-level=info 2
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.


- 2** To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Example output

```
INFO Waiting up to 30m0s for the Kubernetes API at https://api.test.example.com:6443...  
INFO API v1.29.4 up  
INFO Waiting up to 30m0s for bootstrapping to complete...  
INFO It is now safe to remove the bootstrap resources
```

The command succeeds when the Kubernetes API server signals that it has been bootstrapped on the control plane machines.

2. After the bootstrap process is complete, remove the bootstrap machine from the load balancer.



IMPORTANT

You must remove the bootstrap machine from the load balancer at this point. You can also remove or reformat the bootstrap machine itself.

18.7.14. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

18.7.15. Approving the certificate signing requests for your machines

When you add machines to a cluster, two pending certificate signing requests (CSRs) are generated for each machine that you added. You must confirm that these CSRs are approved or, if necessary, approve them yourself. The client requests must be approved first, followed by the server requests.

Prerequisites

- You added machines to your cluster.

Procedure

1. Confirm that the cluster recognizes the machines:

```
$ oc get nodes
```

Example output

```
NAME      STATUS   ROLES    AGE   VERSION
master-0  Ready    master   63m   v1.29.4
master-1  Ready    master   63m   v1.29.4
master-2  Ready    master   64m   v1.29.4
```

The output lists all of the machines that you created.



NOTE

The preceding output might not include the compute nodes, also known as worker nodes, until some CSRs are approved.

2. Review the pending CSRs and ensure that you see the client requests with the **Pending** or **Approved** status for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

NAME	AGE	REQUESTOR	CONDITION
csr-8b2br	15m	system:serviceaccount:openshift-machine-config-operator:node-bootstrapper	Pending
csr-8vnps	15m	system:serviceaccount:openshift-machine-config-operator:node-bootstrapper	Pending
...			

In this example, two machines are joining the cluster. You might see more approved CSRs in the list.

- If the CSRs were not approved, after all of the pending CSRs for the machines you added are in **Pending** status, approve the CSRs for your cluster machines:



NOTE

Because the CSRs rotate automatically, approve your CSRs within an hour of adding the machines to the cluster. If you do not approve them within an hour, the certificates will rotate, and more than two certificates will be present for each node. You must approve all of these certificates. After the client CSR is approved, the Kubelet creates a secondary CSR for the serving certificate, which requires manual approval. Then, subsequent serving certificate renewal requests are automatically approved by the **machine-approver** if the Kubelet requests a new certificate with identical parameters.



NOTE

For clusters running on platforms that are not machine API enabled, such as bare metal and other user-provisioned infrastructure, you must implement a method of automatically approving the kubelet serving certificate requests (CSRs). If a request is not approved, then the **oc exec**, **oc rsh**, and **oc logs** commands cannot succeed, because a serving certificate is required when the API server connects to the kubelet. Any operation that contacts the Kubelet endpoint requires this certificate approval to be in place. The method must watch for new CSRs, confirm that the CSR was submitted by the **node-bootstrapper** service account in the **system:node** or **system:admin** groups, and confirm the identity of the node.

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{"\n"}\n{{end}}{{end}}' | xargs --no-run-if-empty oc adm certificate approve
```



NOTE

Some Operators might not become available until some CSRs are approved.

4. Now that your client requests are approved, you must review the server requests for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

```
NAME      AGE   REQUESTOR                                     CONDITION
csr-bfd72 5m26s system:node:ip-10-0-50-126.us-east-2.compute.internal
Pending
csr-c57lv 5m26s system:node:ip-10-0-95-157.us-east-2.compute.internal
Pending
...
```

5. If the remaining CSRs are not approved, and are in the **Pending** status, approve the CSRs for your cluster machines:

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{\n"}\n{{end}}' | xargs oc adm certificate approve
```

6. After all client and server CSRs have been approved, the machines have the **Ready** status. Verify this by running the following command:

```
$ oc get nodes
```

Example output

```
NAME      STATUS  ROLES  AGE  VERSION
master-0  Ready   master 73m  v1.29.4
master-1  Ready   master 73m  v1.29.4
master-2  Ready   master 74m  v1.29.4
worker-0  Ready   worker 11m  v1.29.4
worker-1  Ready   worker 11m  v1.29.4
```



NOTE

It can take a few minutes after approval of the server CSRs for the machines to transition to the **Ready** status.

Additional information

- For more information on CSRs, see [Certificate Signing Requests](#).

18.7.16. Initial Operator configuration

After the control plane initializes, you must immediately configure some Operators so that they all become available.

Prerequisites

- Your control plane has initialized.

Procedure

1. Watch the cluster components come online:

```
$ watch -n5 oc get clusteroperators
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.16.0	True	False	False	19m
baremetal	4.16.0	True	False	False	37m
cloud-credential	4.16.0	True	False	False	40m
cluster-autoscaler	4.16.0	True	False	False	37m
config-operator	4.16.0	True	False	False	38m
console	4.16.0	True	False	False	26m
csi-snapshot-controller	4.16.0	True	False	False	37m
dns	4.16.0	True	False	False	37m
etcd	4.16.0	True	False	False	36m
image-registry	4.16.0	True	False	False	31m
ingress	4.16.0	True	False	False	30m
insights	4.16.0	True	False	False	31m
kube-apiserver	4.16.0	True	False	False	26m
kube-controller-manager	4.16.0	True	False	False	36m
kube-scheduler	4.16.0	True	False	False	36m
kube-storage-version-migrator	4.16.0	True	False	False	37m
machine-api	4.16.0	True	False	False	29m
machine-approver	4.16.0	True	False	False	37m
machine-config	4.16.0	True	False	False	36m
marketplace	4.16.0	True	False	False	37m
monitoring	4.16.0	True	False	False	29m
network	4.16.0	True	False	False	38m
node-tuning	4.16.0	True	False	False	37m
openshift-apiserver	4.16.0	True	False	False	32m
openshift-controller-manager	4.16.0	True	False	False	30m
openshift-samples	4.16.0	True	False	False	32m
operator-lifecycle-manager	4.16.0	True	False	False	37m
operator-lifecycle-manager-catalog	4.16.0	True	False	False	37m
operator-lifecycle-manager-packageserver	4.16.0	True	False	False	32m
service-ca	4.16.0	True	False	False	38m
storage	4.16.0	True	False	False	37m

2. Configure the Operators that are not available.

18.7.16.1. Disabling the default OperatorHub catalog sources

Operator catalogs that source content provided by Red Hat and community projects are configured for OperatorHub by default during an OpenShift Container Platform installation. In a restricted network environment, you must disable the default catalogs as a cluster administrator.

Procedure

- Disable the sources for the default catalogs by adding **disableAllDefaultSources: true** to the **OperatorHub** object:

```
$ oc patch OperatorHub cluster --type json \
  -p '[{"op": "add", "path": "/spec/disableAllDefaultSources", "value": true}]'
```

TIP

Alternatively, you can use the web console to manage catalog sources. From the **Administration** → **Cluster Settings** → **Configuration** → **OperatorHub** page, click the **Sources** tab, where you can create, update, delete, disable, and enable individual sources.

18.7.16.2. Image registry storage configuration

The Image Registry Operator is not initially available for platforms that do not provide default storage. After installation, you must configure your registry to use storage so that the Registry Operator is made available.

Instructions are shown for configuring a persistent volume, which is required for production clusters. Where applicable, instructions are shown for configuring an empty directory as the storage location, which is available for only non-production clusters.

Additional instructions are provided for allowing the image registry to use block storage types by using the **Recreate** rollout strategy during upgrades.

18.7.16.2.1. Configuring registry storage for IBM Z

As a cluster administrator, following installation you must configure your registry to use storage.

Prerequisites

- You have access to the cluster as a user with the **cluster-admin** role.
- You have a cluster on IBM Z®.
- You have provisioned persistent storage for your cluster, such as Red Hat OpenShift Data Foundation.



IMPORTANT

OpenShift Container Platform supports **ReadWriteOnce** access for image registry storage when you have only one replica. **ReadWriteOnce** access also requires that the registry uses the **Recreate** rollout strategy. To deploy an image registry that supports high availability with two or more replicas, **ReadWriteMany** access is required.

- Must have 100Gi capacity.

Procedure

1. To configure your registry to use storage, change the **spec.storage.pvc** in the **configs.imageregistry/cluster** resource.



NOTE

When you use shared storage, review your security settings to prevent outside access.

2. Verify that you do not have a registry pod:

```
$ oc get pod -n openshift-image-registry -l docker-registry=default
```

Example output

```
No resources found in openshift-image-registry namespace
```



NOTE

If you do have a registry pod in your output, you do not need to continue with this procedure.

3. Check the registry configuration:

```
$ oc edit configs.imageregistry.operator.openshift.io
```

Example output

```
storage:
  pvc:
    claim:
```

Leave the **claim** field blank to allow the automatic creation of an **image-registry-storage** PVC.

4. Check the **clusteroperator** status:

```
$ oc get clusteroperator image-registry
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
image-registry	4.16	True	False	False	6h50m

5. Ensure that your registry is set to managed to enable building and pushing of images.

- Run:

```
$ oc edit configs.imageregistry/cluster
```

Then, change the line


```
managementState: Removed
```

to

```
managementState: Managed
```

18.7.16.2.2. Configuring storage for the image registry in non-production clusters

You must configure storage for the Image Registry Operator. For non-production clusters, you can set the image registry to an empty directory. If you do so, all images are lost if you restart the registry.

Procedure

- To set the image registry storage to an empty directory:

```
$ oc patch configs.imageregistry.operator.openshift.io cluster --type merge --patch '{"spec": {"storage":{"emptyDir":{}}}'
```



WARNING

Configure this option for only non-production clusters.

If you run this command before the Image Registry Operator initializes its components, the **oc patch** command fails with the following error:

```
Error from server (NotFound): configs.imageregistry.operator.openshift.io "cluster" not found
```

Wait a few minutes and run the command again.

18.7.17. Completing installation on user-provisioned infrastructure

After you complete the Operator configuration, you can finish installing the cluster on infrastructure that you provide.

Prerequisites

- Your control plane has initialized.
- You have completed the initial Operator configuration.

Procedure

1. Confirm that all the cluster components are online with the following command:

```
$ watch -n5 oc get clusteroperators
```

Example output

NAME SINCE	VERSION	AVAILABLE	PROGRESSING	DEGRADED	
authentication	4.16.0	True	False	False	19m
baremetal	4.16.0	True	False	False	37m
cloud-credential	4.16.0	True	False	False	40m
cluster-autoscaler	4.16.0	True	False	False	37m
config-operator	4.16.0	True	False	False	38m
console	4.16.0	True	False	False	26m
csi-snapshot-controller	4.16.0	True	False	False	37m
dns	4.16.0	True	False	False	37m
etcd	4.16.0	True	False	False	36m
image-registry	4.16.0	True	False	False	31m
ingress	4.16.0	True	False	False	30m
insights	4.16.0	True	False	False	31m
kube-apiserver	4.16.0	True	False	False	26m
kube-controller-manager	4.16.0	True	False	False	36m
kube-scheduler	4.16.0	True	False	False	36m
kube-storage-version-migrator	4.16.0	True	False	False	37m
machine-api	4.16.0	True	False	False	29m
machine-approver	4.16.0	True	False	False	37m
machine-config	4.16.0	True	False	False	36m
marketplace	4.16.0	True	False	False	37m
monitoring	4.16.0	True	False	False	29m
network	4.16.0	True	False	False	38m
node-tuning	4.16.0	True	False	False	37m
openshift-apiserver	4.16.0	True	False	False	32m
openshift-controller-manager	4.16.0	True	False	False	30m
openshift-samples	4.16.0	True	False	False	32m
operator-lifecycle-manager	4.16.0	True	False	False	37m
operator-lifecycle-manager-catalog	4.16.0	True	False	False	37m
operator-lifecycle-manager-packageserver	4.16.0	True	False	False	32m
service-ca	4.16.0	True	False	False	38m
storage	4.16.0	True	False	False	37m

Alternatively, the following command notifies you when all of the clusters are available. It also retrieves and displays credentials:

```
$ ./openshift-install --dir <installation_directory> wait-for install-complete 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

Example output

```
INFO Waiting up to 30m0s for the cluster to initialize...
```

The command succeeds when the Cluster Version Operator finishes deploying the OpenShift Container Platform cluster from Kubernetes API server.



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

2. Confirm that the Kubernetes API server is communicating with the pods.

a. To view a list of all pods, use the following command:

```
$ oc get pods --all-namespaces
```

Example output

```

NAMESPACE           NAME                                     READY  STATUS
RESTARTS  AGE
openshift-apiserver-operator  openshift-apiserver-operator-85cb746d55-zqhs8  1/1
Running   1    9m
openshift-apiserver          apiserver-67b9g                                1/1  Running  0
3m
openshift-apiserver          apiserver-ljcmx                                1/1  Running  0
1m
openshift-apiserver          apiserver-z25h4                                1/1  Running  0
2m
openshift-authentication-operator  authentication-operator-69d5d8bf84-vh2n8  1/1
Running   0    5m
...

```

b. View the logs for a pod that is listed in the output of the previous command by using the following command:

```
$ oc logs <pod_name> -n <namespace> ❶
```

- ❶ Specify the pod name and namespace, as shown in the output of the previous command.

If the pod logs display, the Kubernetes API server can communicate with the cluster machines.

3. For an installation with Fibre Channel Protocol (FCP), additional steps are required to enable multipathing. Do not enable multipathing during installation. See "Enabling multipathing with kernel arguments on RHCOS" in the *Postinstallation machine configuration tasks* documentation for more information.

4. Register your cluster on the [Cluster registration](#) page.

Verification

If you have enabled secure boot during the OpenShift Container Platform bootstrap process, the following verification steps are required:

1. Debug the node by running the following command:

```
$ oc debug node/<node_name>
chroot /host
```

2. Confirm that secure boot is enabled by running the following command:

```
$ cat /sys/firmware/ipl/secure
```

Example output

```
1 1
```

- 1** The value is **1** if secure boot is enabled and **0** if secure boot is not enabled.

3. List the re-IPL configuration by running the following command:

```
# lsreipl
```

Example output for an FCP disk

```
Re-IPL type: fcp
WWPN: 0x500507630400d1e3
LUN: 0x4001400e00000000
Device: 0.0.810e
bootprog: 0
br_lba: 0
Loadparm: ""
Bootparms: ""
clear: 0
```

Example output for a DASD disk

```
for DASD output:
Re-IPL type: ccw
Device: 0.0.525d
Loadparm: ""
clear: 0
```

4. Shut down the node by running the following command:

```
sudo shutdown -h
```

5. Initiate a boot from LPAR from the Hardware Management Console (HMC). See [Initiating a secure boot from an LPAR](#) in IBM documentation.

- When the node is back, check the secure boot status again.

Additional resources

- [How to generate SOSREPORT within OpenShift Container Platform version 4 nodes without SSH.](#)

18.7.18. Next steps

- [Customize your cluster.](#)
- If the mirror registry that you used to install your cluster has a trusted CA, add it to the cluster by [configuring additional trust stores](#).
- If necessary, you can [opt out of remote health reporting](#).
- If necessary, see [Registering your disconnected cluster](#)

18.8. INSTALLATION CONFIGURATION PARAMETERS FOR IBM Z AND IBM LINUXONE

Before you deploy an OpenShift Container Platform cluster, you provide a customized **install-config.yaml** installation configuration file that describes the details for your environment.



NOTE

While this document refers only to IBM Z[®], all information in it also applies to IBM[®] LinuxONE.

18.8.1. Available installation configuration parameters for IBM Z

The following tables specify the required, optional, and IBM Z-specific installation configuration parameters that you can set as part of the installation process.



NOTE

After installation, you cannot modify these parameters in the **install-config.yaml** file.

18.8.1.1. Required configuration parameters

Required installation configuration parameters are described in the following table:

Table 18.114. Required parameters

Parameter	Description	Values
apiVersion:	The API version for the install-config.yaml content. The current version is v1 . The installation program may also support older API versions.	String

Parameter	Description	Values
<code>baseDomain:</code>	The base domain of your cloud provider. The base domain is used to create routes to your OpenShift Container Platform cluster components. The full DNS name for your cluster is a combination of the baseDomain and metadata.name parameter values that uses the <metadata.name> . <baseDomain> format.	A fully-qualified domain or subdomain name, such as example.com .
<code>metadata:</code>	Kubernetes resource ObjectMeta , from which only the name parameter is consumed.	Object
<code>metadata: name:</code>	The name of the cluster. DNS records for the cluster are all subdomains of {{.metadata.name}} . {{.baseDomain}} .	String of lowercase letters, hyphens (-), and periods (.), such as dev .
<code>platform:</code>	The configuration for the specific platform upon which to perform the installation: aws, baremetal, azure, gcp, ibmcloud, nutanix, openstack, powervs, vsphere , or {} . For additional information about platform . <platform> parameters, consult the table for your specific platform that follows.	Object
<code>pullSecret:</code>	Get a pull secret from Red Hat OpenShift Cluster Manager to authenticate downloading container images for OpenShift Container Platform components from services such as Quay.io.	<pre>{ "auths":{ "cloud.openshift.com":{ "auth":"b3Blb=", "email":"you@example.com" }, "quay.io":{ "auth":"b3Blb=", "email":"you@example.com" } } }</pre>

18.8.1.2. Network configuration parameters

You can customize your installation configuration based on the requirements of your existing network infrastructure. For example, you can expand the IP address block for the cluster network or provide different IP address blocks than the defaults.

- If you use the Red Hat OpenShift Networking OVN-Kubernetes network plugin, both IPv4 and IPv6 address families are supported.

If you configure your cluster to use both IP address families, review the following requirements:

- Both IP families must use the same network interface for the default gateway.
- Both IP families must have the default gateway.
- You must specify IPv4 and IPv6 addresses in the same order for all network configuration parameters. For example, in the following configuration IPv4 addresses are listed before IPv6 addresses.


```
networking:
  clusterNetwork:
    - cidr: 10.128.0.0/14
      hostPrefix: 23
    - cidr: fd00:10:128::/56
      hostPrefix: 64
  serviceNetwork:
    - 172.30.0.0/16
    - fd00:172:16::/112
```




NOTE

Globalnet is not supported with Red Hat OpenShift Data Foundation disaster recovery solutions. For regional disaster recovery scenarios, ensure that you use a nonoverlapping range of private IP addresses for the cluster and service networks in each cluster.

Table 18.115. Network parameters

Parameter	Description	Values
networking:	The configuration for the cluster network.	Object  NOTE You cannot modify parameters specified by the networking object after installation.

Parameter	Description	Values
<code>networking: networkType:</code>	The Red Hat OpenShift Networking network plugin to install.	OVNKubernetes. OVNKubernetes is a CNI plugin for Linux networks and hybrid networks that contain both Linux and Windows servers. The default value is OVNKubernetes .
<code>networking: clusterNetwork:</code>	The IP address blocks for pods. The default value is 10.128.0.0/14 with a host prefix of /23 . If you specify multiple IP address blocks, the blocks must not overlap.	An array of objects. For example: <code>networking: clusterNetwork: - cidr: 10.128.0.0/14 hostPrefix: 23</code>
<code>networking: clusterNetwork: cidr:</code>	Required if you use networking.clusterNetwork . An IP address block. An IPv4 network.	An IP address block in Classless Inter-Domain Routing (CIDR) notation. The prefix length for an IPv4 block is between 0 and 32 .
<code>networking: clusterNetwork: hostPrefix:</code>	The subnet prefix length to assign to each individual node. For example, if hostPrefix is set to 23 then each node is assigned a /23 subnet out of the given cidr . A hostPrefix value of 23 provides $2^{(32 - 23) - 2}$ pod IP addresses.	A subnet prefix. The default value is 23 .
<code>networking: serviceNetwork:</code>	The IP address block for services. The default value is 172.30.0.0/16 . The OVN-Kubernetes network plugins supports only a single IP address block for the service network.	An array with an IP address block in CIDR format. For example: <code>networking: serviceNetwork: - 172.30.0.0/16</code>
<code>networking: machineNetwork:</code>	The IP address blocks for machines. If you specify multiple IP address blocks, the blocks must not overlap. If you specify multiple IP kernel arguments, the machineNetwork.cidr value must be the CIDR of the primary network.	An array of objects. For example: <code>networking: machineNetwork: - cidr: 10.0.0.0/16</code>


Parameter	Description	Values
networking: machineNetwork: cidr:	Required if you use networking.machineNetwork . An IP address block. The default value is 10.0.0.0/16 for all platforms other than libvirt and IBM Power® Virtual Server. For libvirt, the default value is 192.168.126.0/24 . For IBM Power® Virtual Server, the default value is 192.168.0.0/24 .	An IP network block in CIDR notation. For example, 10.0.0.0/16 .  <div style="margin-left: 20px;"> NOTE Set the networking.machineNetwork to match the CIDR that the preferred NIC resides in. </div>


18.8.1.3. Optional configuration parameters

Optional installation configuration parameters are described in the following table:

Table 18.116. Optional parameters

Parameter	Description	Values
additionalTrustBundle:	A PEM-encoded X.509 certificate bundle that is added to the nodes' trusted certificate store. This trust bundle may also be used when a proxy has been configured.	String
capabilities:	Controls the installation of optional core cluster components. You can reduce the footprint of your OpenShift Container Platform cluster by disabling optional components. For more information, see the "Cluster capabilities" page in <i>Installing</i> .	String array
capabilities: baselineCapabilitySet:	Selects an initial set of optional capabilities to enable. Valid values are None , v4.11 , v4.12 and vCurrent . The default value is vCurrent .	String
capabilities: additionalEnabledCapabilities:	Extends the set of optional capabilities beyond what you specify in baselineCapabilitySet . You may specify multiple capabilities in this parameter.	String array

Parameter	Description	Values
<code>cpuPartitioningMode:</code>	Enables workload partitioning, which isolates OpenShift Container Platform services, cluster management workloads, and infrastructure pods to run on a reserved set of CPUs. Workload partitioning can only be enabled during installation and cannot be disabled after installation. While this field enables workload partitioning, it does not configure workloads to use specific CPUs. For more information, see the <i>Workload partitioning</i> page in the <i>Scalability and Performance</i> section.	None or AllNodes . None is the default value.
<code>compute:</code>	The configuration for the machines that comprise the compute nodes.	Array of MachinePool objects.
<code>compute: architecture:</code>	Determines the instruction set architecture of the machines in the pool. Currently, heterogeneous clusters are not supported, so all pools must specify the same architecture. Valid values are s390x (the default).	String
<code>compute: hyperthreading:</code>	Whether to enable or disable simultaneous multithreading, or hyperthreading , on compute machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.  IMPORTANT If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.	Enabled or Disabled
<code>compute: name:</code>	Required if you use compute . The name of the machine pool.	worker

Parameter	Description	Values
<code>compute: platform:</code>	Required if you use compute . Use this parameter to specify the cloud provider to host the worker machines. This parameter value must match the controlPlane.platform parameter value.	aws, azure, gcp, ibmcloud, nutanix, openstack, powervs, vsphere , or <code>{}</code>
<code>compute: replicas:</code>	The number of compute machines, which are also known as worker machines, to provision.	A positive integer greater than or equal to 2 . The default value is 3 .
<code>featureSet:</code>	Enables the cluster for a feature set. A feature set is a collection of OpenShift Container Platform features that are not enabled by default. For more information about enabling a feature set during installation, see "Enabling features using feature gates".	String. The name of the feature set to enable, such as TechPreviewNoUpgrade .
<code>controlPlane:</code>	The configuration for the machines that comprise the control plane.	Array of MachinePool objects.
<code>controlPlane: architecture:</code>	Determines the instruction set architecture of the machines in the pool. Currently, heterogeneous clusters are not supported, so all pools must specify the same architecture. Valid values are s390x (the default).	String
<code>controlPlane: hyperthreading:</code>	Whether to enable or disable simultaneous multithreading, or hyperthreading , on control plane machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.  IMPORTANT If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.	Enabled or Disabled

Parameter	Description	Values
<code>controlPlane: name:</code>	Required if you use controlPlane . The name of the machine pool.	master
<code>controlPlane: platform:</code>	Required if you use controlPlane . Use this parameter to specify the cloud provider that hosts the control plane machines. This parameter value must match the compute.platform parameter value.	aws, azure, gcp, ibmcloud, nutanix, openstack, powervs, vsphere , or {}
<code>controlPlane: replicas:</code>	The number of control plane machines to provision.	Supported values are 3 , or 1 when deploying single-node OpenShift.
<code>credentialsMode:</code>	The Cloud Credential Operator (CCO) mode. If no mode is specified, the CCO dynamically tries to determine the capabilities of the provided credentials, with a preference for mint mode on the platforms where multiple modes are supported.	Mint, Passthrough, Manual or an empty string (<code>""</code>). ^[1]

Parameter	Description	Values
<p>fips:</p>	<p>Enable or disable FIPS mode. The default is false (disabled). If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.</p> <div data-bbox="486 584 592 1637" style="background-color: black; color: white; padding: 5px;"> <p>IMPORTANT</p> <p>To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see Installing the system in FIPS mode.</p> <p>When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.</p> </div> <div data-bbox="486 1682 592 1883" style="background-color: #f0f0f0; padding: 5px;"> <p>NOTE</p> <p>If you are using Azure File storage, you cannot enable FIPS mode.</p> </div>	<p>false or true</p>

Parameter	Description	Values
<code>imageContentSources:</code>	Sources and repositories for the release-image content.	Array of objects. Includes a source and, optionally, mirrors , as described in the following rows of this table.
<code>imageContentSources: source:</code>	Required if you use imageContentSources . Specify the repository that users refer to, for example, in image pull specifications.	String
<code>imageContentSources: mirrors:</code>	Specify one or more repositories that may also contain the same images.	Array of strings
<code>publish:</code>	How to publish or expose the user-facing endpoints of your cluster, such as the Kubernetes API, OpenShift routes.	<p>Internal or External. The default value is External.</p> <p>Setting this field to Internal is not supported on non-cloud platforms.</p> <div style="display: flex; align-items: flex-start;"> <div style="background-color: #ccc; width: 20px; height: 20px; margin-right: 5px;"></div> <div> <p>IMPORTANT</p> <p>If the value of the field is set to Internal, the cluster will become non-functional. For more information, refer to BZ#1953035.</p> </div> </div>
<code>sshKey:</code>	<p>The SSH key to authenticate access to your cluster machines.</p> <div style="display: flex; align-items: flex-start;"> <div style="background-color: #ccc; width: 20px; height: 20px; margin-right: 5px;"></div> <div> <p>NOTE</p> <p>For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your ssh-agent process uses.</p> </div> </div>	For example, sshKey: ssh-ed25519 AAAA...

1. Not all CCO modes are supported for all cloud providers. For more information about CCO modes, see the "Managing cloud provider credentials" entry in the *Authentication and authorization* content.

CHAPTER 19. INSTALLING ON IBM POWER

19.1. PREPARING TO INSTALL ON IBM POWER

19.1.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).


19.1.2. Choosing a method to install OpenShift Container Platform on IBM Power

You can install a cluster on IBM Power® infrastructure that you provision, by using one of the following methods:

- **Installing a cluster on IBM Power®**: You can install OpenShift Container Platform on IBM Power® infrastructure that you provision.
- **Installing a cluster on IBM Power® in a restricted network**: You can install OpenShift Container Platform on IBM Power® infrastructure that you provision in a restricted or disconnected network, by using an internal mirror of the installation release content. You can use this method to install a cluster that does not require an active internet connection to obtain the software components. You can also use this installation method to ensure that your clusters only use container images that satisfy your organizational controls on external content.

19.2. INSTALLING A CLUSTER ON IBM POWER

In OpenShift Container Platform version 4.16, you can install a cluster on IBM Power® infrastructure that you provision.



IMPORTANT

Additional considerations exist for non-bare metal platforms. Review the information in the [guidelines for deploying OpenShift Container Platform on non-tested platforms](#) before you install an OpenShift Container Platform cluster.

19.2.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- Before you begin the installation process, you must clean the installation directory. This ensures that the required installation files are created and updated during the installation process.
- You provisioned [persistent storage using OpenShift Data Foundation](#) or other supported storage protocols for your cluster. To deploy a private image registry, you must set up persistent storage with **ReadWriteMany** access.

- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.

**NOTE**


Be sure to also review this site list if you are configuring a proxy.

19.2.2. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.

**IMPORTANT**

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

19.2.3. Requirements for a cluster with user-provisioned infrastructure

For a cluster that contains user-provisioned infrastructure, you must deploy all of the required machines.

This section describes the requirements for deploying OpenShift Container Platform on user-provisioned infrastructure.


19.2.3.1. Required machines for cluster installation

The smallest OpenShift Container Platform clusters require the following hosts:

Table 19.1. Minimum required hosts

Hosts	Description
One temporary bootstrap machine	The cluster requires the bootstrap machine to deploy the OpenShift Container Platform cluster on the three control plane machines. You can remove the bootstrap machine after you install the cluster.
Three control plane machines	The control plane machines run the Kubernetes and OpenShift Container Platform services that form the control plane.

Hosts	Description
At least two compute machines, which are also known as worker machines.	The workloads requested by OpenShift Container Platform users run on the compute machines.



IMPORTANT

To maintain high availability of your cluster, use separate physical hosts for these cluster machines.

The bootstrap, control plane, and compute machines must use Red Hat Enterprise Linux CoreOS (RHCOS) as the operating system.

Note that RHCOS is based on Red Hat Enterprise Linux (RHEL) 9.2 and inherits all of its hardware certifications and requirements. See [Red Hat Enterprise Linux technology capabilities and limits](#) .

19.2.3.2. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 19.2. Minimum resource requirements

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Bootstrap	RHCOS	2	16 GB	100 GB	300
Control plane	RHCOS	2	16 GB	100 GB	300
Compute	RHCOS	2	8 GB	100 GB	300

1. One vCPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio: $(\text{threads per core} \times \text{cores}) \times \text{sockets} = \text{vCPUs}$.
2. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes. Note that on many cloud platforms, storage size and IOPS scale together, so you might need to over-allocate storage volume to obtain sufficient performance.

**NOTE**

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

19.2.3.3. Minimum IBM Power requirements

You can install OpenShift Container Platform version 4.16 on the following IBM® hardware:

- IBM Power®9 or IBM Power®10 processor-based systems

**NOTE**

Support for RHCOS functionality for all IBM Power®8 models, IBM Power® AC922, IBM Power® IC922, and IBM Power® LC922 is deprecated in OpenShift Container Platform 4.16. Red Hat recommends that you use later hardware models.

Hardware requirements

- Six logical partitions (LPARs) across multiple PowerVM servers

Operating system requirements

- One instance of an IBM Power®9 or Power10 processor-based system

On your IBM Power® instance, set up:

- Three LPARs for OpenShift Container Platform control plane machines
- Two LPARs for OpenShift Container Platform compute machines
- One LPAR for the temporary OpenShift Container Platform bootstrap machine

Disk storage for the IBM Power guest virtual machines

- Local storage, or storage provisioned by the Virtual I/O Server using vSCSI, NPIV (N-Port ID Virtualization) or SSP (shared storage pools)

Network for the PowerVM guest virtual machines

- Dedicated physical adapter, or SR-IOV virtual function
- Available by the Virtual I/O Server using Shared Ethernet Adapter
- Virtualized by the Virtual I/O Server using IBM® vNIC

Storage / main memory

- 100 GB / 16 GB for OpenShift Container Platform control plane machines
- 100 GB / 8 GB for OpenShift Container Platform compute machines
- 100 GB / 16 GB for the temporary OpenShift Container Platform bootstrap machine

19.2.3.4. Recommended IBM Power system requirements**Hardware requirements**

- Six LPARs across multiple PowerVM servers

Operating system requirements

- One instance of an IBM Power®9 or IBM Power®10 processor-based system

On your IBM Power® instance, set up:

- Three LPARs for OpenShift Container Platform control plane machines
- Two LPARs for OpenShift Container Platform compute machines
- One LPAR for the temporary OpenShift Container Platform bootstrap machine

Disk storage for the IBM Power guest virtual machines

- Local storage, or storage provisioned by the Virtual I/O Server using vSCSI, NPIV (N-Port ID Virtualization) or SSP (shared storage pools)

Network for the PowerVM guest virtual machines

- Dedicated physical adapter, or SR-IOV virtual function
- Virtualized by the Virtual I/O Server using Shared Ethernet Adapter
- Virtualized by the Virtual I/O Server using IBM® vNIC

Storage / main memory

- 120 GB / 32 GB for OpenShift Container Platform control plane machines
- 120 GB / 32 GB for OpenShift Container Platform compute machines
- 120 GB / 16 GB for the temporary OpenShift Container Platform bootstrap machine

19.2.3.5. Certificate signing requests management

Because your cluster has limited access to automatic machine management when you use infrastructure

that you provision, you must provide a mechanism for approving cluster certificate signing requests (CSRs) after installation. The **kube-controller-manager** only approves the kubelet client CSRs. The **machine-approver** cannot guarantee the validity of a serving certificate that is requested by using kubelet credentials because it cannot confirm that the correct machine issued the request. You must determine and implement a method of verifying the validity of the kubelet serving certificate requests and approving them.

19.2.3.6. Networking requirements for user-provisioned infrastructure

All the Red Hat Enterprise Linux CoreOS (RHCOS) machines require networking to be configured in **initramfs** during boot to fetch their Ignition config files.

During the initial boot, the machines require an IP address configuration that is set either through a DHCP server or statically by providing the required boot options. After a network connection is established, the machines download their Ignition config files from an HTTP or HTTPS server. The Ignition config files are then used to set the exact state of each machine. The Machine Config Operator completes more changes to the machines, such as the application of new certificates or keys, after installation.

It is recommended to use a DHCP server for long-term management of the cluster machines. Ensure that the DHCP server is configured to provide persistent IP addresses, DNS server information, and hostnames to the cluster machines.



NOTE

If a DHCP service is not available for your user-provisioned infrastructure, you can instead provide the IP networking configuration and the address of the DNS server to the nodes at RHCOS install time. These can be passed as boot arguments if you are installing from an ISO image. See the *Installing RHCOS and starting the OpenShift Container Platform bootstrap process* section for more information about static IP provisioning and advanced networking options.

The Kubernetes API server must be able to resolve the node names of the cluster machines. If the API servers and worker nodes are in different zones, you can configure a default DNS search zone to allow the API server to resolve the node names. Another supported approach is to always refer to hosts by their fully-qualified domain names in both the node objects and all DNS requests.

19.2.3.6.1. Setting the cluster node hostnames through DHCP

On Red Hat Enterprise Linux CoreOS (RHCOS) machines, the hostname is set through NetworkManager. By default, the machines obtain their hostname through DHCP. If the hostname is not provided by DHCP, set statically through kernel arguments, or another method, it is obtained through a reverse DNS lookup. Reverse DNS lookup occurs after the network has been initialized on a node and can take time to resolve. Other system services can start prior to this and detect the hostname as **localhost** or similar. You can avoid this by using DHCP to provide the hostname for each cluster node.

Additionally, setting the hostnames through DHCP can bypass any manual DNS record name configuration errors in environments that have a DNS split-horizon implementation.

19.2.3.6.2. Network connectivity requirements

You must configure the network connectivity between machines to allow OpenShift Container Platform cluster components to communicate. Each machine must be able to resolve the hostnames of all other machines in the cluster.

This section provides details about the ports that are required.



IMPORTANT

In connected OpenShift Container Platform environments, all nodes are required to have internet access to pull images for platform containers and provide telemetry data to Red Hat.

Table 19.3. Ports used for all-machine to all-machine communications

Protocol	Port	Description
ICMP	N/A	Network reachability tests
TCP	1936	Metrics
	9000-9999	Host level services, including the node exporter on ports 9100-9101 and the Cluster Version Operator on port 9099 .
	10250-10259	The default ports that Kubernetes reserves
UDP	4789	VXLAN
	6081	Geneve
	9000-9999	Host level services, including the node exporter on ports 9100-9101 .
	500	IPsec IKE packets
	4500	IPsec NAT-T packets
	123	Network Time Protocol (NTP) on UDP port 123 If an external NTP time server is configured, you must open UDP port 123 .
TCP/UDP	30000-32767	Kubernetes node port
ESP	N/A	IPsec Encapsulating Security Payload (ESP)

Table 19.4. Ports used for all-machine to control plane communications

Protocol	Port	Description
TCP	6443	Kubernetes API

Table 19.5. Ports used for control plane machine to control plane machine communications

Protocol	Port	Description
TCP	2379-2380	etcd server and peer ports

NTP configuration for user-provisioned infrastructure

OpenShift Container Platform clusters are configured to use a public Network Time Protocol (NTP) server by default. If you want to use a local enterprise NTP server, or if your cluster is being deployed in a disconnected network, you can configure the cluster to use a specific time server. For more information, see the documentation for *Configuring chrony time service*.

If a DHCP server provides NTP server information, the chrony time service on the Red Hat Enterprise Linux CoreOS (RHCOS) machines read the information and can sync the clock with the NTP servers.

Additional resources

- [Configuring chrony time service](#)

19.2.3.7. User-provisioned DNS requirements

In OpenShift Container Platform deployments, DNS name resolution is required for the following components:

- The Kubernetes API
- The OpenShift Container Platform application wildcard
- The bootstrap, control plane, and compute machines

Reverse DNS resolution is also required for the Kubernetes API, the bootstrap machine, the control plane machines, and the compute machines.

DNS A/AAAA or CNAME records are used for name resolution and PTR records are used for reverse name resolution. The reverse records are important because Red Hat Enterprise Linux CoreOS (RHCOS) uses the reverse records to set the hostnames for all the nodes, unless the hostnames are provided by DHCP. Additionally, the reverse records are used to generate the certificate signing requests (CSR) that OpenShift Container Platform needs to operate.





NOTE

It is recommended to use a DHCP server to provide the hostnames to each cluster node. See the *DHCP recommendations for user-provisioned infrastructure* section for more information.

The following DNS records are required for a user-provisioned OpenShift Container Platform cluster and they must be in place before installation. In each record, **<cluster_name>** is the cluster name and **<base_domain>** is the base domain that you specify in the **install-config.yaml** file. A complete DNS record takes the form: **<component>.<cluster_name>.<base_domain>..**

Table 19.6. Required DNS records

Component	Record	Description
Kubernetes API	api.<cluster_name>.<base_domain>.	A DNS A/AAAA or CNAME record, and a DNS PTR record, to identify the API load balancer. These records must be resolvable by both clients external to the cluster and from all the nodes within the cluster.
	api-int.<cluster_name>.<base_domain>.	<p>A DNS A/AAAA or CNAME record, and a DNS PTR record, to internally identify the API load balancer. These records must be resolvable from all the nodes within the cluster.</p> <div style="display: flex; align-items: flex-start;">  <div> <p>IMPORTANT</p> <p>The API server must be able to resolve the worker nodes by the hostnames that are recorded in Kubernetes. If the API server cannot resolve the node names, then proxied API calls can fail, and you cannot retrieve logs from pods.</p> </div> </div>
Routes	*.apps.<cluster_name>.<base_domain>.	<p>A wildcard DNS A/AAAA or CNAME record that refers to the application ingress load balancer. The application ingress load balancer targets the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default. These records must be resolvable by both clients external to the cluster and from all the nodes within the cluster.</p> <p>For example, console-openshift-console.apps.<cluster_name>.<base_domain> is used as a wildcard route to the OpenShift Container Platform console.</p>
Bootstrap machine	bootstrap.<cluster_name>.<base_domain>.	A DNS A/AAAA or CNAME record, and a DNS PTR record, to identify the bootstrap machine. These records must be resolvable by the nodes within the cluster.
Control plane machines	<control_plane><n>.<cluster_name>.<base_domain>.	DNS A/AAAA or CNAME records and DNS PTR records to identify each machine for the control plane nodes. These records must be resolvable by the nodes within the cluster.
Compute machines	<compute><n>.<cluster_name>.<base_domain>.	DNS A/AAAA or CNAME records and DNS PTR records to identify each machine for the worker nodes. These records must be resolvable by the nodes within the cluster.

**NOTE**

In OpenShift Container Platform 4.4 and later, you do not need to specify etcd host and SRV records in your DNS configuration.

TIP

You can use the **dig** command to verify name and reverse name resolution. See the section on *Validating DNS resolution for user-provisioned infrastructure* for detailed validation steps.

19.2.3.7.1. Example DNS configuration for user-provisioned clusters

This section provides A and PTR record configuration samples that meet the DNS requirements for deploying OpenShift Container Platform on user-provisioned infrastructure. The samples are not meant to provide advice for choosing one DNS solution over another.

In the examples, the cluster name is **ocp4** and the base domain is **example.com**.

Example DNS A record configuration for a user-provisioned cluster

The following example is a BIND zone file that shows sample A records for name resolution in a user-provisioned cluster.

Example 19.1. Sample DNS zone database

```
$TTL 1W
@ IN SOA ns1.example.com. root (
    2019070700 ; serial
    3H ; refresh (3 hours)
    30M ; retry (30 minutes)
    2W ; expiry (2 weeks)
    1W ) ; minimum (1 week)
IN NS ns1.example.com.
IN MX 10 smtp.example.com.
;
;
ns1.example.com. IN A 192.168.1.5
smtp.example.com. IN A 192.168.1.5
;
helper.example.com. IN A 192.168.1.5
helper.ocp4.example.com. IN A 192.168.1.5
;
api.ocp4.example.com. IN A 192.168.1.5 1
api-int.ocp4.example.com. IN A 192.168.1.5 2
;
*.apps.ocp4.example.com. IN A 192.168.1.5 3
;
bootstrap.ocp4.example.com. IN A 192.168.1.96 4
;
control-plane0.ocp4.example.com. IN A 192.168.1.97 5
control-plane1.ocp4.example.com. IN A 192.168.1.98 6
control-plane2.ocp4.example.com. IN A 192.168.1.99 7
;
compute0.ocp4.example.com. IN A 192.168.1.11 8
```



```
compute1.ocp4.example.com. IN A 192.168.1.7 9
```

```
;
;EOF
```

- 1 Provides name resolution for the Kubernetes API. The record refers to the IP address of the API load balancer.
- 2 Provides name resolution for the Kubernetes API. The record refers to the IP address of the API load balancer and is used for internal cluster communications.
- 3 Provides name resolution for the wildcard routes. The record refers to the IP address of the application ingress load balancer. The application ingress load balancer targets the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.



NOTE

In the example, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

- 4 Provides name resolution for the bootstrap machine.
- 5 6 7 Provides name resolution for the control plane machines.
- 8 9 Provides name resolution for the compute machines.

Example DNS PTR record configuration for a user-provisioned cluster

The following example BIND zone file shows sample PTR records for reverse name resolution in a user-provisioned cluster.

Example 19.2. Sample DNS zone database for reverse records

```
$TTL 1W
@ IN SOA ns1.example.com. root (
  2019070700 ; serial
  3H ; refresh (3 hours)
  30M ; retry (30 minutes)
  2W ; expiry (2 weeks)
  1W ) ; minimum (1 week)
IN NS ns1.example.com.
;
5.1.168.192.in-addr.arpa. IN PTR api.ocp4.example.com. 1
5.1.168.192.in-addr.arpa. IN PTR api-int.ocp4.example.com. 2
;
96.1.168.192.in-addr.arpa. IN PTR bootstrap.ocp4.example.com. 3
;
97.1.168.192.in-addr.arpa. IN PTR control-plane0.ocp4.example.com. 4
98.1.168.192.in-addr.arpa. IN PTR control-plane1.ocp4.example.com. 5
99.1.168.192.in-addr.arpa. IN PTR control-plane2.ocp4.example.com. 6
```

```

;
11.1.168.192.in-addr.arpa. IN PTR compute0.ocp4.example.com. 7
7.1.168.192.in-addr.arpa. IN PTR compute1.ocp4.example.com. 8
;
;EOF

```

- 1 Provides reverse DNS resolution for the Kubernetes API. The PTR record refers to the record name of the API load balancer.
- 2 Provides reverse DNS resolution for the Kubernetes API. The PTR record refers to the record name of the API load balancer and is used for internal cluster communications.
- 3 Provides reverse DNS resolution for the bootstrap machine.
- 4 5 6 Provides reverse DNS resolution for the control plane machines.
- 7 8 Provides reverse DNS resolution for the compute machines.

**NOTE**

A PTR record is not required for the OpenShift Container Platform application wildcard.

19.2.3.8. Load balancing requirements for user-provisioned infrastructure


Before you install OpenShift Container Platform, you must provision the API and application Ingress load balancing infrastructure. In production scenarios, you can deploy the API and application Ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

**NOTE**

If you want to deploy the API and application Ingress load balancers with a Red Hat Enterprise Linux (RHEL) instance, you must purchase the RHEL subscription separately.

The load balancing infrastructure must meet the following requirements:

1. **API load balancer.** Provides a common endpoint for users, both human and machine, to interact with and configure the platform. Configure the following conditions:
 - Layer 4 load balancing only. This can be referred to as Raw TCP or SSL Passthrough mode.
 - A stateless load balancing algorithm. The options vary based on the load balancer implementation.

**IMPORTANT**

Do not configure session persistence for an API load balancer. Configuring session persistence for a Kubernetes API server might cause performance issues from excess application traffic for your OpenShift Container Platform cluster and the Kubernetes API that runs inside the cluster.

Configure the following ports on both the front and back of the load balancers:

Table 19.7. API load balancer

Port	Back-end machines (pool members)	Internal	External	Description
6443	Bootstrap and control plane. You remove the bootstrap machine from the load balancer after the bootstrap machine initializes the cluster control plane. You must configure the /readyz endpoint for the API server health check probe.	X	X	Kubernetes API server
22623	Bootstrap and control plane. You remove the bootstrap machine from the load balancer after the bootstrap machine initializes the cluster control plane.	X		Machine config server

**NOTE**

The load balancer must be configured to take a maximum of 30 seconds from the time the API server turns off the **/readyz** endpoint to the removal of the API server instance from the pool. Within the time frame after **/readyz** returns an error or becomes healthy, the endpoint must have been removed or added. Probing every 5 or 10 seconds, with two successful requests to become healthy and three to become unhealthy, are well-tested values.

2. **Application Ingress load balancer:** Provides an ingress point for application traffic flowing in from outside the cluster. A working configuration for the Ingress router is required for an OpenShift Container Platform cluster.

Configure the following conditions:

- Layer 4 load balancing only. This can be referred to as Raw TCP or SSL Passthrough mode.
- A connection-based or session-based persistence is recommended, based on the options available and types of applications that will be hosted on the platform.

TIP

If the true IP address of the client can be seen by the application Ingress load balancer, enabling source IP-based session persistence can improve performance for applications that use end-to-end TLS encryption.

Configure the following ports on both the front and back of the load balancers:

Table 19.8. Application Ingress load balancer

Port	Back-end machines (pool members)	Internal	External	Description
443	The machines that run the Ingress Controller pods, compute, or worker, by default.	X	X	HTTPS traffic
80	The machines that run the Ingress Controller pods, compute, or worker, by default.	X	X	HTTP traffic

**NOTE**

If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application Ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes.

19.2.3.8.1. Example load balancer configuration for user-provisioned clusters

This section provides an example API and application Ingress load balancer configuration that meets the load balancing requirements for user-provisioned clusters. The sample is an `/etc/haproxy/haproxy.cfg` configuration for an HAProxy load balancer. The example is not meant to provide advice for choosing one load balancing solution over another.

In the example, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

**NOTE**

If you are using HAProxy as a load balancer and SELinux is set to **enforcing**, you must ensure that the HAProxy service can bind to the configured TCP port by running **setsebool -P haproxy_connect_any=1**.

Example 19.3. Sample API and application Ingress load balancer configuration


```
global
  log 127.0.0.1 local2
  pidfile /var/run/haproxy.pid
  maxconn 4000
  daemon
defaults
  mode http
  log global
  option dontlognull
  option http-server-close
  option redispatch
  retries 3
  timeout http-request 10s
  timeout queue 1m
  timeout connect 10s
  timeout client 1m
```

```

timeout server      1m
timeout http-keep-alive 10s
timeout check      10s
maxconn            3000
listen api-server-6443 1
bind *:6443
mode tcp
option httpchk GET /readyz HTTP/1.0
option log-health-checks
balance roundrobin
server bootstrap bootstrap.ocp4.example.com:6443 verify none check check-ssl inter 10s fall 2
rise 3 backup 2
server master0 master0.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
fall 2 rise 3
server master1 master1.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
fall 2 rise 3
server master2 master2.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
fall 2 rise 3
listen machine-config-server-22623 3
bind *:22623
mode tcp
server bootstrap bootstrap.ocp4.example.com:22623 check inter 1s backup 4
server master0 master0.ocp4.example.com:22623 check inter 1s
server master1 master1.ocp4.example.com:22623 check inter 1s
server master2 master2.ocp4.example.com:22623 check inter 1s
listen ingress-router-443 5
bind *:443
mode tcp
balance source
server compute0 compute0.ocp4.example.com:443 check inter 1s
server compute1 compute1.ocp4.example.com:443 check inter 1s
listen ingress-router-80 6
bind *:80
mode tcp
balance source
server compute0 compute0.ocp4.example.com:80 check inter 1s
server compute1 compute1.ocp4.example.com:80 check inter 1s

```

- 1** Port **6443** handles the Kubernetes API traffic and points to the control plane machines.
- 2** **4** The bootstrap entries must be in place before the OpenShift Container Platform cluster installation and they must be removed after the bootstrap process is complete.
- 3** Port **22623** handles the machine config server traffic and points to the control plane machines.
- 5** Port **443** handles the HTTPS traffic and points to the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.
- 6** Port **80** handles the HTTP traffic and points to the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.

**NOTE**

If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application Ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes.

TIP

If you are using HAProxy as a load balancer, you can check that the **haproxy** process is listening on ports **6443**, **22623**, **443**, and **80** by running **netstat -nltp** on the HAProxy node.

19.2.4. Preparing the user-provisioned infrastructure

Before you install OpenShift Container Platform on user-provisioned infrastructure, you must prepare the underlying infrastructure.

This section provides details about the high-level steps required to set up your cluster infrastructure in preparation for an OpenShift Container Platform installation. This includes configuring IP networking and network connectivity for your cluster nodes, enabling the required ports through your firewall, and setting up the required DNS and load balancing infrastructure.

After preparation, your cluster infrastructure must meet the requirements outlined in the *Requirements for a cluster with user-provisioned infrastructure* section.

Prerequisites

- You have reviewed the [OpenShift Container Platform 4.x Tested Integrations](#) page.
- You have reviewed the infrastructure requirements detailed in the *Requirements for a cluster with user-provisioned infrastructure* section.

Procedure

1. If you are using DHCP to provide the IP networking configuration to your cluster nodes, configure your DHCP service.
 - a. Add persistent IP addresses for the nodes to your DHCP server configuration. In your configuration, match the MAC address of the relevant network interface to the intended IP address for each node.
 - b. When you use DHCP to configure IP addressing for the cluster machines, the machines also obtain the DNS server information through DHCP. Define the persistent DNS server address that is used by the cluster nodes through your DHCP server configuration.

**NOTE**

If you are not using a DHCP service, you must provide the IP networking configuration and the address of the DNS server to the nodes at RHCOS install time. These can be passed as boot arguments if you are installing from an ISO image. See the *Installing RHCOS and starting the OpenShift Container Platform bootstrap process* section for more information about static IP provisioning and advanced networking options.

- c. Define the hostnames of your cluster nodes in your DHCP server configuration. See the *Setting the cluster node hostnames through DHCP* section for details about hostname considerations.

**NOTE**

If you are not using a DHCP service, the cluster nodes obtain their hostname through a reverse DNS lookup.

2. Ensure that your network infrastructure provides the required network connectivity between the cluster components. See the *Networking requirements for user-provisioned infrastructure* section for details about the requirements.
3. Configure your firewall to enable the ports required for the OpenShift Container Platform cluster components to communicate. See *Networking requirements for user-provisioned infrastructure* section for details about the ports that are required.

**IMPORTANT**

By default, port **1936** is accessible for an OpenShift Container Platform cluster, because each control plane node needs access to this port.

Avoid using the Ingress load balancer to expose this port, because doing so might result in the exposure of sensitive information, such as statistics and metrics, related to Ingress Controllers.


4. Setup the required DNS infrastructure for your cluster.
 - a. Configure DNS name resolution for the Kubernetes API, the application wildcard, the bootstrap machine, the control plane machines, and the compute machines.
 - b. Configure reverse DNS resolution for the Kubernetes API, the bootstrap machine, the control plane machines, and the compute machines.
See the *User-provisioned DNS requirements* section for more information about the OpenShift Container Platform DNS requirements.
5. Validate your DNS configuration.
 - a. From your installation node, run DNS lookups against the record names of the Kubernetes API, the wildcard routes, and the cluster nodes. Validate that the IP addresses in the responses correspond to the correct components.
 - b. From your installation node, run reverse DNS lookups against the IP addresses of the load balancer and the cluster nodes. Validate that the record names in the responses correspond to the correct components.
See the *Validating DNS resolution for user-provisioned infrastructure* section for detailed DNS validation steps.
6. Provision the required API and application ingress load balancing infrastructure. See the *Load balancing requirements for user-provisioned infrastructure* section for more information about the requirements.

**NOTE**

Some load balancing solutions require the DNS name resolution for the cluster nodes to be in place before the load balancing is initialized.

19.2.5. Validating DNS resolution for user-provisioned infrastructure

You can validate your DNS configuration before installing OpenShift Container Platform on user-provisioned infrastructure.

**IMPORTANT**

The validation steps detailed in this section must succeed before you install your cluster.

Prerequisites

- You have configured the required DNS records for your user-provisioned infrastructure.

Procedure

- From your installation node, run DNS lookups against the record names of the Kubernetes API, the wildcard routes, and the cluster nodes. Validate that the IP addresses contained in the responses correspond to the correct components.
 - Perform a lookup against the Kubernetes API record name. Check that the result points to the IP address of the API load balancer:

```
$ dig +noall +answer @<nameserver_ip> api.<cluster_name>.<base_domain> 1
```

- Replace **<nameserver_ip>** with the IP address of the nameserver, **<cluster_name>** with your cluster name, and **<base_domain>** with your base domain name.

Example output

```
api.ocp4.example.com. 604800 IN A 192.168.1.5
```

- Perform a lookup against the Kubernetes internal API record name. Check that the result points to the IP address of the API load balancer:

```
$ dig +noall +answer @<nameserver_ip> api-int.<cluster_name>.<base_domain>
```

Example output

```
api-int.ocp4.example.com. 604800 IN A 192.168.1.5
```

- Test an example ***.apps.<cluster_name>.<base_domain>** DNS wildcard lookup. All of the application wildcard lookups must resolve to the IP address of the application ingress load balancer:

```
$ dig +noall +answer @<nameserver_ip> random.apps.<cluster_name>.<base_domain>
```

Example output


```
random.apps.ocp4.example.com. 604800 IN A 192.168.1.5
```



NOTE

In the example outputs, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

You can replace **random** with another wildcard value. For example, you can query the route to the OpenShift Container Platform console:

```
$ dig +noall +answer @<nameserver_ip> console-openshift-console.apps.
<cluster_name>.<base_domain>
```

Example output

```
console-openshift-console.apps.ocp4.example.com. 604800 IN A 192.168.1.5
```

- d. Run a lookup against the bootstrap DNS record name. Check that the result points to the IP address of the bootstrap node:

```
$ dig +noall +answer @<nameserver_ip> bootstrap.<cluster_name>.<base_domain>
```

Example output

```
bootstrap.ocp4.example.com. 604800 IN A 192.168.1.96
```

- e. Use this method to perform lookups against the DNS record names for the control plane and compute nodes. Check that the results correspond to the IP addresses of each node.
2. From your installation node, run reverse DNS lookups against the IP addresses of the load balancer and the cluster nodes. Validate that the record names contained in the responses correspond to the correct components.
 - a. Perform a reverse lookup against the IP address of the API load balancer. Check that the response includes the record names for the Kubernetes API and the Kubernetes internal API:

```
$ dig +noall +answer @<nameserver_ip> -x 192.168.1.5
```

Example output

```
5.1.168.192.in-addr.arpa. 604800 IN PTR api-int.ocp4.example.com. 1
5.1.168.192.in-addr.arpa. 604800 IN PTR api.ocp4.example.com. 2
```

1 Provides the record name for the Kubernetes internal API.

2 Provides the record name for the Kubernetes API.

**NOTE**

A PTR record is not required for the OpenShift Container Platform application wildcard. No validation step is needed for reverse DNS resolution against the IP address of the application ingress load balancer.

- b. Perform a reverse lookup against the IP address of the bootstrap node. Check that the result points to the DNS record name of the bootstrap node:

```
$ dig +noall +answer @<nameserver_ip> -x 192.168.1.96
```

Example output

```
96.1.168.192.in-addr.arpa. 604800 IN PTR bootstrap.ocp4.example.com.
```

- c. Use this method to perform reverse lookups against the IP addresses for the control plane and compute nodes. Check that the results correspond to the DNS record names of each node.

19.2.6. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.

**IMPORTANT**

Do not skip this procedure in production environments, where disaster recovery and debugging is required.

**NOTE**


You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.



NOTE

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.


2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.



NOTE

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```



NOTE

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1 Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
-
```

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

19.2.7. Obtaining the installation program

Before you install OpenShift Container Platform, download the installation file on the host you are using for installation.

Prerequisites

- You have a computer that runs Linux or macOS, with 500 MB of local disk space.

Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select your infrastructure provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.



IMPORTANT

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.



IMPORTANT

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.

4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#) . This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

19.2.8. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.



IMPORTANT

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**.
To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

■

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:


```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.



NOTE

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your PATH.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

19.2.9. Manually creating the installation configuration file

Installing the cluster requires that you manually create the installation configuration file.


Prerequisites

- You have an SSH public key on your local machine to provide to the installation program. The key will be used for SSH authentication onto your cluster nodes for debugging and disaster recovery.
- You have obtained the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create an installation directory to store your required installation assets in:

```
$ mkdir <installation_directory>
```



IMPORTANT

You must create a directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.


2. Customize the sample **install-config.yaml** file template that is provided and save it in the **<installation_directory>**.



NOTE

You must name this configuration file **install-config.yaml**.

3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the next step of the installation process. You must back it up now.

Additional resources

- [Installation configuration parameters for IBM Power®](#)

19.2.9.1. Sample install-config.yaml file for IBM Power

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.

```
apiVersion: v1
baseDomain: example.com 1
compute: 2
- hyperthreading: Enabled 3
  name: worker
  replicas: 0 4
  architecture: ppc64le
controlPlane: 5
  hyperthreading: Enabled 6
  name: master
  replicas: 3 7
  architecture: ppc64le
metadata:
  name: test 8
networking:
  clusterNetwork:
```

```

- cidr: 10.128.0.0/14 9
  hostPrefix: 23 10
  networkType: OVNKubernetes 11
  serviceNetwork: 12
  - 172.30.0.0/16
platform:
  none: {} 13
fips: false 14
pullSecret: '{"auths": ...}' 15
sshKey: 'ssh-ed25519 AAAA...' 16

```

- ¹ The base domain of the cluster. All DNS records must be sub-domains of this base and include the cluster name.
- ² ⁵ The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.
- ³ ⁶ Simultaneous multithreading (SMT) is not supported.
- ⁴ You must set this value to **0** when you install OpenShift Container Platform on user-provisioned infrastructure. In installer-provisioned installations, the parameter controls the number of compute machines that the cluster creates and manages for you. In user-provisioned installations, you must manually deploy the compute machines before you finish installing the cluster.



NOTE

If you are installing a three-node cluster, do not deploy any compute machines when you install the Red Hat Enterprise Linux CoreOS (RHCOS) machines.

- ⁷ The number of control plane machines that you add to the cluster. Because the cluster uses these values as the number of etcd endpoints in the cluster, the value must match the number of control plane machines that you deploy.
- ⁸ The cluster name that you specified in your DNS records.
- ⁹ A block of IP addresses from which pod IP addresses are allocated. This block must not overlap with existing physical networks. These IP addresses are used for the pod network. If you need to access the pods from an external network, you must configure load balancers and routers to manage the traffic.



NOTE

Class E CIDR range is reserved for a future use. To use the Class E CIDR range, you must ensure your networking environment accepts the IP addresses within the Class E CIDR range.

- ¹⁰ The subnet prefix length to assign to each individual node. For example, if **hostPrefix** is set to **23**, then each node is assigned a **/23** subnet out of the given **cidr**, which allows for 510 ($2^{(32 - 23)} - 2$) pod IP addresses. If you are required to provide access to nodes from an external network, configure load balancers and routers to manage the traffic.
- ¹¹ The cluster network plugin to install. The default value **OVNKubernetes** is the only supported

value.

- 12 The IP address pool to use for service IP addresses. You can enter only one IP address pool. This block must not overlap with existing physical networks. If you need to access the services from an external network, configure load balancers and routers to manage the traffic.
- 13 You must set the platform to **none**. You cannot provide additional platform configuration variables for IBM Power® infrastructure.



IMPORTANT

Clusters that are installed with the platform type **none** are unable to use some features, such as managing compute machines with the Machine API. This limitation applies even if the compute machines that are attached to the cluster are installed on a platform that would normally support the feature. This parameter cannot be changed after installation.

- 14 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.




IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 15 The [pull secret from Red Hat OpenShift Cluster Manager](#) . This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.
- 16 The SSH public key for the **core** user in Red Hat Enterprise Linux CoreOS (RHCOS).



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

19.2.9.2. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).


Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxvonly** and

`user-ca-bundle` config map in the `credentials` field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the `user-ca-bundle` config map only when `http/https` proxy is configured. Use **Always** to always reference the `user-ca-bundle` config map. The default value is **Proxyonly**.



NOTE

The installation program does not support the proxy `readinessEndpoints` field.



NOTE

If the installer times out, restart and then complete the deployment by using the `wait-for` command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named `cluster` that uses the proxy settings in the provided `install-config.yaml` file. If no proxy settings are provided, a `cluster Proxy` object is still created, but it will have a nil `spec`.



NOTE

Only the `Proxy` object named `cluster` is supported, and no additional proxies can be created.

19.2.9.3. Configuring a three-node cluster

Optionally, you can deploy zero compute machines in a bare metal cluster that consists of three control plane machines only. This provides smaller, more resource efficient clusters for cluster administrators and developers to use for testing, development, and production.

In three-node OpenShift Container Platform environments, the three control plane machines are schedulable, which means that your application workloads are scheduled to run on them.

Prerequisites

- You have an existing `install-config.yaml` file.

Procedure

- Ensure that the number of compute replicas is set to `0` in your `install-config.yaml` file, as shown in the following `compute` stanza:

```
compute:
  - name: worker
    platform: {}
    replicas: 0
```

**NOTE**

You must set the value of the **replicas** parameter for the compute machines to **0** when you install OpenShift Container Platform on user-provisioned infrastructure, regardless of the number of compute machines you are deploying. In installer-provisioned installations, the parameter controls the number of compute machines that the cluster creates and manages for you. This does not apply to user-provisioned installations, where the compute machines are deployed manually.

For three-node cluster installations, follow these next steps:

- If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes. See the *Load balancing requirements for user-provisioned infrastructure* section for more information.
- When you create the Kubernetes manifest files in the following procedure, ensure that the **mastersSchedulable** parameter in the `<installation_directory>/manifests/cluster-scheduler-02-config.yml` file is set to **true**. This enables your application workloads to run on the control plane nodes.
- Do not deploy any compute nodes when you create the Red Hat Enterprise Linux CoreOS (RHCOS) machines.

19.2.10. Cluster Network Operator configuration

The configuration for the cluster network is specified as part of the Cluster Network Operator (CNO) configuration and stored in a custom resource (CR) object that is named **cluster**. The CR specifies the fields for the **Network** API in the **operator.openshift.io** API group.

The CNO configuration inherits the following fields during cluster installation from the **Network** API in the **Network.config.openshift.io** API group:

clusterNetwork

IP address pools from which pod IP addresses are allocated.

serviceNetwork

IP address pool for services.

defaultNetwork.type

Cluster network plugin. **OVNKubernetes** is the only supported plugin during installation.

You can specify the cluster network plugin configuration for your cluster by setting the fields for the **defaultNetwork** object in the CNO object named **cluster**.

19.2.10.1. Cluster Network Operator configuration object

The fields for the Cluster Network Operator (CNO) are described in the following table:

Table 19.9. Cluster Network Operator configuration object


Field	Type	Description
metadata.name	string	The name of the CNO object. This name is always cluster .
spec.clusterNetwork	array	<p>A list specifying the blocks of IP addresses from which pod IP addresses are allocated and the subnet prefix length assigned to each individual node in the cluster. For example:</p> <pre>spec: clusterNetwork: - cidr: 10.128.0.0/19 hostPrefix: 23 - cidr: 10.128.32.0/19 hostPrefix: 23</pre>
spec.serviceNetwork	array	<p>A block of IP addresses for services. The OpenShift SDN and OVN-Kubernetes network plugins support only a single IP address block for the service network. For example:</p> <pre>spec: serviceNetwork: - 172.30.0.0/14</pre> <p>You can customize this field only in the install-config.yaml file before you create the manifests. The value is read-only in the manifest file.</p>
spec.defaultNetwork	object	Configures the network plugin for the cluster network.
spec.kubeProxyConfig	object	The fields for this object specify the kube-proxy configuration. If you are using the OVN-Kubernetes cluster network plugin, the kube-proxy configuration has no effect.

defaultNetwork object configuration

The values for the **defaultNetwork** object are defined in the following table:

Table 19.10. **defaultNetwork** object

Field	Type	Description
-------	------	-------------

Field	Type	Description
type	string	<p>OVNKubernetes. The Red Hat OpenShift Networking network plugin is selected during installation. This value cannot be changed after cluster installation.</p> <div style="display: flex; align-items: flex-start;">  <div> <p>NOTE</p> <p>OpenShift Container Platform uses the OVN-Kubernetes network plugin by default. OpenShift SDN is no longer available as an installation choice for new clusters.</p> </div> </div>
ovnKubernetesConfig	object	This object is only valid for the OVN-Kubernetes network plugin.

Configuration for the OVN-Kubernetes network plugin

The following table describes the configuration fields for the OVN-Kubernetes network plugin:

Table 19.11. **ovnKubernetesConfig** object

Field	Type	Description
mtu	integer	<p>The maximum transmission unit (MTU) for the Geneve (Generic Network Virtualization Encapsulation) overlay network. This is detected automatically based on the MTU of the primary network interface. You do not normally need to override the detected MTU.</p> <p>If the auto-detected value is not what you expect it to be, confirm that the MTU on the primary network interface on your nodes is correct. You cannot use this option to change the MTU value of the primary network interface on the nodes.</p> <p>If your cluster requires different MTU values for different nodes, you must set this value to 100 less than the lowest MTU value in your cluster. For example, if some nodes in your cluster have an MTU of 9001, and some have an MTU of 1500, you must set this value to 1400.</p>
genevePort	integer	The port to use for all Geneve packets. The default value is 6081 . This value cannot be changed after cluster installation.
ipsecConfig	object	Specify a configuration object for customizing the IPsec configuration.
ipv4	object	Specifies a configuration object for IPv4 settings.
ipv6	object	Specifies a configuration object for IPv6 settings.


Field	Type	Description
policyAuditConfig	object	Specify a configuration object for customizing network policy audit logging. If unset, the defaults audit log settings are used.
gatewayConfig	object	Optional: Specify a configuration object for customizing how egress traffic is sent to the node gateway.  <p>NOTE</p> <p>While migrating egress traffic, you can expect some disruption to workloads and service traffic until the Cluster Network Operator (CNO) successfully rolls out the changes.</p>

Table 19.12. `ovnKubernetesConfig.ipv4` object

Field	Type	Description
internalTransitSwitchSubnet	string	If your existing network infrastructure overlaps with the 100.88.0.0/16 IPv4 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. The subnet for the distributed transit switch that enables east-west traffic. This subnet cannot overlap with any other subnets used by OVN-Kubernetes or on the host itself. It must be large enough to accommodate one IP address per node in your cluster. The default value is 100.88.0.0/16 .
internalJoinSubnet	string	If your existing network infrastructure overlaps with the 100.64.0.0/16 IPv4 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. You must ensure that the IP address range does not overlap with any other subnet used by your OpenShift Container Platform installation. The IP address range must be larger than the maximum number of nodes that can be added to the cluster. For example, if the clusterNetwork.cidr value is 10.128.0.0/14 and the clusterNetwork.hostPrefix value is /23 , then the maximum number of nodes is 2⁽²³⁻¹⁴⁾=512 . The default value is 100.64.0.0/16 .

Table 19.13. `ovnKubernetesConfig.ipv6` object

Field	Type	Description
-------	------	-------------

Field	Type	Description
internalTransitSwitchSubnet	string	<p>If your existing network infrastructure overlaps with the fd97::/64 IPv6 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. The subnet for the distributed transit switch that enables east-west traffic. This subnet cannot overlap with any other subnets used by OVN-Kubernetes or on the host itself. It must be large enough to accommodate one IP address per node in your cluster.</p> <p>The default value is fd97::/64.</p>
internalJoinSubnet	string	<p>If your existing network infrastructure overlaps with the fd98::/64 IPv6 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. You must ensure that the IP address range does not overlap with any other subnet used by your OpenShift Container Platform installation. The IP address range must be larger than the maximum number of nodes that can be added to the cluster.</p> <p>The default value is fd98::/64.</p>

Table 19.14. `policyAuditConfig` object

Field	Type	Description
rateLimit	integer	The maximum number of messages to generate every second per node. The default value is 20 messages per second.
maxFileSize	integer	The maximum size for the audit log in bytes. The default value is 50000000 or 50 MB.
maxLogFiles	integer	The maximum number of log files that are retained.
destination	string	<p>One of the following additional audit log targets:</p> <p>libc The libc syslog() function of the journald process on the host.</p> <p>udp:<host>:<port> A syslog server. Replace <host>:<port> with the host and port of the syslog server.</p> <p>unix:<file> A Unix Domain Socket file specified by <file>.</p> <p>null Do not send the audit logs to any additional target.</p>
syslogFacility	string	The syslog facility, such as kern , as defined by RFC5424. The default value is local0 .

Table 19.15. gatewayConfig object

Field	Type	Description
routingViaHost	boolean	<p>Set this field to true to send egress traffic from pods to the host networking stack. For highly-specialized installations and applications that rely on manually configured routes in the kernel routing table, you might want to route egress traffic to the host networking stack. By default, egress traffic is processed in OVN to exit the cluster and is not affected by specialized routes in the kernel routing table. The default value is false.</p> <p>This field has an interaction with the Open vSwitch hardware offloading feature. If you set this field to true, you do not receive the performance benefits of the offloading because egress traffic is processed by the host networking stack.</p>
ipForwarding	object	<p>You can control IP forwarding for all traffic on OVN-Kubernetes managed interfaces by using the ipForwarding specification in the Network resource. Specify Restricted to only allow IP forwarding for Kubernetes related traffic. Specify Global to allow forwarding of all IP traffic. For new installations, the default is Restricted. For updates to OpenShift Container Platform 4.14 or later, the default is Global.</p>
ipv4	object	<p>Optional: Specify an object to configure the internal OVN-Kubernetes masquerade address for host to service traffic for IPv4 addresses.</p>
ipv6	object	<p>Optional: Specify an object to configure the internal OVN-Kubernetes masquerade address for host to service traffic for IPv6 addresses.</p>

Table 19.16. gatewayConfig.ipv4 object

Field	Type	Description
internalMasqueradeSubnet	string	<p>The masquerade IPv4 addresses that are used internally to enable host to service traffic. The host is configured with these IP addresses as well as the shared gateway bridge interface. The default value is 169.254.169.0/29.</p>

Table 19.17. gatewayConfig.ipv6 object

Field	Type	Description
internalMasqueradeSubnet	string	<p>The masquerade IPv6 addresses that are used internally to enable host to service traffic. The host is configured with these IP addresses as well as the shared gateway bridge interface. The default value is fd69::/125.</p>

Table 19.18. ipsecConfig object

Field	Type	Description
mode	string	<p>Specifies the behavior of the IPsec implementation. Must be one of the following values:</p> <ul style="list-style-type: none"> ● Disabled: IPsec is not enabled on cluster nodes. ● External: IPsec is enabled for network traffic with external hosts. ● Full: IPsec is enabled for pod traffic and network traffic with external hosts.

Example OVN-Kubernetes configuration with IPsec enabled

```
defaultNetwork:
  type: OVNKubernetes
  ovnKubernetesConfig:
    mtu: 1400
    genevePort: 6081
    ipsecConfig:
      mode: Full
```




IMPORTANT

Using OVNKubernetes can lead to a stack exhaustion problem on IBM Power®.

kubeProxyConfig object configuration (OpenShiftSDN container network interface only)

The values for the **kubeProxyConfig** object are defined in the following table:

Table 19.19. kubeProxyConfig object


Field	Type	Description
iptablesSyncPeriod	string	<p>The refresh period for iptables rules. The default value is 30s. Valid suffixes include s, m, and h and are described in the Go time package documentation.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p>NOTE</p> <p>Because of performance improvements introduced in OpenShift Container Platform 4.3 and greater, adjusting the iptablesSyncPeriod parameter is no longer necessary.</p> </div> </div>

Field	Type	Description
proxyArguments.iptables-min-sync-period	array	<p>The minimum duration before refreshing iptables rules. This field ensures that the refresh does not happen too frequently. Valid suffixes include s, m, and h and are described in the Go time package. The default value is:</p> <pre>kubeProxyConfig: proxyArguments: iptables-min-sync-period: - 0s</pre>

19.2.11. Creating the Kubernetes manifest and Ignition config files

Because you must modify some cluster definition files and manually start the cluster machines, you must generate the Kubernetes manifest and Ignition config files that the cluster needs to configure the machines.

The installation configuration file transforms into the Kubernetes manifests. The manifests wrap into the Ignition configuration files, which are later used to configure the cluster machines.



IMPORTANT

- The Ignition config files that the OpenShift Container Platform installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.



NOTE

The installation program that generates the manifest and Ignition files is architecture specific and can be obtained from the [client image mirror](#). The Linux version of the installation program (without an architecture postfix) runs on ppc64le only. This installer program is also available as a Mac OS version.

Prerequisites

- You obtained the OpenShift Container Platform installation program.
- You created the **install-config.yaml** installation configuration file.

Procedure

1. Change to the directory that contains the OpenShift Container Platform installation program and generate the Kubernetes manifests for the cluster:


```
$ ./openshift-install create manifests --dir <installation_directory> 1
```

- 1 For **<installation_directory>**, specify the installation directory that contains the **install-config.yaml** file you created.



WARNING

If you are installing a three-node cluster, skip the following step to allow the control plane nodes to be schedulable.



IMPORTANT

When you configure control plane nodes from the default unschedulable to schedulable, additional subscriptions are required. This is because control plane nodes then become compute nodes.

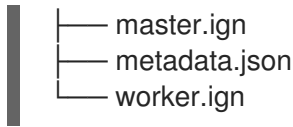
2. Check that the **mastersSchedulable** parameter in the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** Kubernetes manifest file is set to **false**. This setting prevents pods from being scheduled on the control plane machines:
 - a. Open the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** file.
 - b. Locate the **mastersSchedulable** parameter and ensure that it is set to **false**.
 - c. Save and exit the file.
3. To create the Ignition configuration files, run the following command from the directory that contains the installation program:

```
$ ./openshift-install create ignition-configs --dir <installation_directory> 1
```

- 1 For **<installation_directory>**, specify the same installation directory.

Ignition config files are created for the bootstrap, control plane, and compute nodes in the installation directory. The **kubeadmin-password** and **kubeconfig** files are created in the **./<installation_directory>/auth** directory:

```
.
├── auth
│   ├── kubeadmin-password
│   └── kubeconfig
└── bootstrap.ign
```



19.2.12. Installing RHCOS and starting the OpenShift Container Platform bootstrap process

To install OpenShift Container Platform on IBM Power® infrastructure that you provision, you must install Red Hat Enterprise Linux CoreOS (RHCOS) on the machines. When you install RHCOS, you must provide the Ignition config file that was generated by the OpenShift Container Platform installation program for the type of machine you are installing. If you have configured suitable networking, DNS, and load balancing infrastructure, the OpenShift Container Platform bootstrap process begins automatically after the RHCOS machines have rebooted.

Follow either the steps to use an ISO image or network PXE booting to install RHCOS on the machines.

19.2.12.1. Installing RHCOS by using an ISO image

You can use an ISO image to install RHCOS on the machines.

Prerequisites

- You have created the Ignition config files for your cluster.
- You have configured suitable network, DNS and load balancing infrastructure.
- You have an HTTP server that can be accessed from your computer, and from the machines that you create.
- You have reviewed the *Advanced RHCOS installation configuration* section for different ways to configure features, such as networking and disk partitioning.

Procedure

1. Obtain the SHA512 digest for each of your Ignition config files. For example, you can use the following on a system running Linux to get the SHA512 digest for your **bootstrap.ign** Ignition config file:

```
$ sha512sum <installation_directory>/bootstrap.ign
```

The digests are provided to the **coreos-installer** in a later step to validate the authenticity of the Ignition config files on the cluster nodes.

2. Upload the bootstrap, control plane, and compute node Ignition config files that the installation program created to your HTTP server. Note the URLs of these files.



IMPORTANT

You can add or change configuration settings in your Ignition configs before saving them to your HTTP server. If you plan to add more compute machines to your cluster after you finish installation, do not delete these files.

3. From the installation host, validate that the Ignition config files are available on the URLs. The following example gets the Ignition config file for the bootstrap node:

```
$ curl -k http://<HTTP_server>/bootstrap.ign 1
```

Example output

```
% Total % Received % Xferd Average Speed Time Time Time Current
           Dload Upload Total Spent Left Speed
0 0 0 0 0 0 0 0 --:--:-- --:--:-- --:--:-- 0{"ignition":
{"version":"3.2.0"},"passwd":{"users":[{"name":"core","sshAuthorizedKeys":["ssh-rsa...
```

Replace **bootstrap.ign** with **master.ign** or **worker.ign** in the command to validate that the Ignition config files for the control plane and compute nodes are also available.

- Although it is possible to obtain the RHCOS images that are required for your preferred method of installing operating system instances from the [RHCOS image mirror](#) page, the recommended way to obtain the correct version of your RHCOS images are from the output of **openshift-install** command:

```
$ openshift-install coreos print-stream-json | grep '\.iso[^\.]'
```

Example output

```
"location": "<url>/art/storage/releases/rhcos-4.16-aarch64/<release>/aarch64/rhcos-
<release>-live.aarch64.iso",
"location": "<url>/art/storage/releases/rhcos-4.16-ppc64le/<release>/ppc64le/rhcos-
<release>-live.ppc64le.iso",
"location": "<url>/art/storage/releases/rhcos-4.16-s390x/<release>/s390x/rhcos-<release>-
live.s390x.iso",
"location": "<url>/art/storage/releases/rhcos-4.16/<release>/x86_64/rhcos-<release>-
live.x86_64.iso",
```



IMPORTANT

The RHCOS images might not change with every release of OpenShift Container Platform. You must download images with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Use the image versions that match your OpenShift Container Platform version if they are available. Use only ISO images for this procedure. RHCOS qcow2 images are not supported for this installation type.

ISO file names resemble the following example:

rhcos-<version>-live.<architecture>.iso

- Use the ISO to start the RHCOS installation. Use one of the following installation options:
 - Burn the ISO image to a disk and boot it directly.
 - Use ISO redirection by using a lights-out management (LOM) interface.
- Boot the RHCOS ISO image without specifying any options or interrupting the live boot sequence. Wait for the installer to boot into a shell prompt in the RHCOS live environment.

**NOTE**

It is possible to interrupt the RHCOS installation boot process to add kernel arguments. However, for this ISO procedure you should use the **coreos-installer** command as outlined in the following steps, instead of adding kernel arguments.

7. Run the **coreos-installer** command and specify the options that meet your installation requirements. At a minimum, you must specify the URL that points to the Ignition config file for the node type, and the device that you are installing to:

```
$ sudo coreos-installer install --ignition-url=http://<HTTP_server>/<node_type>.ign <device>
--ignition-hash=sha512-<digest> 1 2
```

- 1 1 You must run the **coreos-installer** command by using **sudo**, because the **core** user does not have the required root privileges to perform the installation.
- 2 The **--ignition-hash** option is required when the Ignition config file is obtained through an HTTP URL to validate the authenticity of the Ignition config file on the cluster node. **<digest>** is the Ignition config file SHA512 digest obtained in a preceding step.


**NOTE**

If you want to provide your Ignition config files through an HTTPS server that uses TLS, you can add the internal certificate authority (CA) to the system trust store before running **coreos-installer**.

The following example initializes a bootstrap node installation to the **/dev/sda** device. The Ignition config file for the bootstrap node is obtained from an HTTP web server with the IP address 192.168.1.2:

```
$ sudo coreos-installer install --ignition-
url=http://192.168.1.2:80/installation_directory/bootstrap.ign /dev/sda --ignition-hash=sha512-
a5a2d43879223273c9b60af66b44202a1d1248fc01cf156c46d4a79f552b6bad47bc8cc78ddf011
6e80c59d2ea9e32ba53bc807afbca581aa059311def2c3e3b
```

8. Monitor the progress of the RHCOS installation on the console of the machine.

**IMPORTANT**


Be sure that the installation is successful on each node before commencing with the OpenShift Container Platform installation. Observing the installation process can also help to determine the cause of RHCOS installation issues that might arise.

9. After RHCOS installs, you must reboot the system. During the system reboot, it applies the Ignition config file that you specified.
10. Check the console output to verify that Ignition ran.

Example command

Ignition: ran on 2022/03/14 14:48:33 UTC (this boot)
 Ignition: user-provided config was applied

11. Continue to create the other machines for your cluster.



IMPORTANT

You must create the bootstrap and control plane machines at this time. If the control plane machines are not made schedulable, also create at least two compute machines before you install OpenShift Container Platform.

If the required network, DNS, and load balancer infrastructure are in place, the OpenShift Container Platform bootstrap process begins automatically after the RHCOS nodes have rebooted.



NOTE

RHCOS nodes do not include a default password for the **core** user. You can access the nodes by running **ssh core@<node>.<cluster_name>.<base_domain>** as a user with access to the SSH private key that is paired to the public key that you specified in your **install_config.yaml** file. OpenShift Container Platform 4 cluster nodes running RHCOS are immutable and rely on Operators to apply cluster changes. Accessing cluster nodes by using SSH is not recommended. However, when investigating installation issues, if the OpenShift Container Platform API is not available, or the kubelet is not properly functioning on a target node, SSH access might be required for debugging or disaster recovery.

19.2.12.1.1. Advanced RHCOS installation reference

This section illustrates the networking configuration and other advanced options that allow you to modify the Red Hat Enterprise Linux CoreOS (RHCOS) manual installation process. The following tables describe the kernel arguments and command-line options you can use with the RHCOS live installer and the **coreos-installer** command.

19.2.12.1.1.1. Networking and bonding options for ISO installations


If you install RHCOS from an ISO image, you can add kernel arguments manually when you boot the image to configure networking for a node. If no networking arguments are specified, DHCP is activated in the initramfs when RHCOS detects that networking is required to fetch the Ignition config file.



IMPORTANT

When adding networking arguments manually, you must also add the **rd.neednet=1** kernel argument to bring the network up in the initramfs.

The following information provides examples for configuring networking and bonding on your RHCOS nodes for ISO installations. The examples describe how to use the **ip=**, **nameserver=**, and **bond=** kernel arguments.

**NOTE**

Ordering is important when adding the kernel arguments: **ip=**, **nameserver=**, and then **bond=**.

The networking options are passed to the **dracut** tool during system boot. For more information about the networking options supported by **dracut**, see the [dracut.cmdline manual page](#).


The following examples are the networking options for ISO installation.

Configuring DHCP or static IP addresses

To configure an IP address, either use DHCP (**ip=dhcp**) or set an individual static IP address (**ip=<host_ip>**). If setting a static IP, you must then identify the DNS server IP address (**nameserver=<dns_ip>**) on each node. The following example sets:

- The node's IP address to **10.10.10.2**
- The gateway address to **10.10.10.254**
- The netmask to **255.255.255.0**
- The hostname to **core0.example.com**
- The DNS server address to **4.4.4.41**
- The auto-configuration value to **none**. No auto-configuration is required when IP networking is configured statically.

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none
nameserver=4.4.4.41
```

**NOTE**

When you use DHCP to configure IP addressing for the RHCOS machines, the machines also obtain the DNS server information through DHCP. For DHCP-based deployments, you can define the DNS server address that is used by the RHCOS nodes through your DHCP server configuration.

Configuring an IP address without a static hostname

You can configure an IP address without assigning a static hostname. If a static hostname is not set by the user, it will be picked up and automatically set by a reverse DNS lookup. To configure an IP address without a static hostname refer to the following example:

- The node's IP address to **10.10.10.2**
- The gateway address to **10.10.10.254**
- The netmask to **255.255.255.0**
- The DNS server address to **4.4.4.41**
- The auto-configuration value to **none**. No auto-configuration is required when IP networking is configured statically.

```
ip=10.10.10.2::10.10.10.254:255.255.255.0::enp1s0:none
nameserver=4.4.4.41
```

Specifying multiple network interfaces

You can specify multiple network interfaces by setting multiple **ip=** entries.

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none
ip=10.10.10.3::10.10.10.254:255.255.255.0:core0.example.com:enp2s0:none
```

Configuring default gateway and route

Optional: You can configure routes to additional networks by setting an **rd.route=** value.



NOTE

When you configure one or multiple networks, one default gateway is required. If the additional network gateway is different from the primary network gateway, the default gateway must be the primary network gateway.

- Run the following command to configure the default gateway:

```
ip=::10.10.10.254:::
```

- Enter the following command to configure the route for the additional network:

```
rd.route=20.20.20.0/24:20.20.20.254:enp2s0
```

Disabling DHCP on a single interface

You can disable DHCP on a single interface, such as when there are two or more network interfaces and only one interface is being used. In the example, the **enp1s0** interface has a static networking configuration and DHCP is disabled for **enp2s0**, which is not used:

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none
ip=::::core0.example.com:enp2s0:none
```

Combining DHCP and static IP configurations

You can combine DHCP and static IP configurations on systems with multiple network interfaces, for example:

```
ip=enp1s0:dhcp
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp2s0:none
```

Configuring VLANs on individual interfaces

Optional: You can configure VLANs on individual interfaces by using the **vlan=** parameter.

- To configure a VLAN on a network interface and use a static IP address, run the following command:

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp2s0.100:none
vlan=enp2s0.100:enp2s0
```

- To configure a VLAN on a network interface and to use DHCP, run the following command:

```
ip=enp2s0.100:dhcp
vlan=enp2s0.100:enp2s0
```

Providing multiple DNS servers

You can provide multiple DNS servers by adding a **nameserver=** entry for each server, for example:

```
nameserver=1.1.1.1
nameserver=8.8.8.8
```

Bonding multiple network interfaces to a single interface

Optional: You can bond multiple network interfaces to a single interface by using the **bond=** option. Refer to the following examples:


- The syntax for configuring a bonded interface is: **bond=<name>[:<network_interfaces>][:options]**
<name> is the bonding device name (**bond0**), **<network_interfaces>** represents a comma-separated list of physical (ethernet) interfaces (**em1,em2**), and *options* is a comma-separated list of bonding options. Enter **modinfo bonding** to see available options.
- When you create a bonded interface using **bond=**, you must specify how the IP address is assigned and other information for the bonded interface.
 - To configure the bonded interface to use DHCP, set the bond's IP address to **dhcp**. For example:

```
bond=bond0:em1,em2:mode=active-backup
ip=bond0:dhcp
```

- To configure the bonded interface to use a static IP address, enter the specific IP address you want and related information. For example:

```
bond=bond0:em1,em2:mode=active-backup
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:bond0:none
```

Bonding multiple SR-IOV network interfaces to a dual port NIC interface



IMPORTANT

Support for Day 1 operations associated with enabling NIC partitioning for SR-IOV devices is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

Optional: You can bond multiple SR-IOV network interfaces to a dual port NIC interface by using the **bond=** option.

On each node, you must perform the following tasks:

1. Create the SR-IOV virtual functions (VFs) following the guidance in [Managing SR-IOV devices](#). Follow the procedure in the "Attaching SR-IOV networking devices to virtual machines" section.
2. Create the bond, attach the desired VFs to the bond and set the bond link state up following the guidance in [Configuring network bonding](#). Follow any of the described procedures to create the bond.

The following examples illustrate the syntax you must use:

- The syntax for configuring a bonded interface is **bond=<name>[:<network_interfaces>] [:options]**.
<name> is the bonding device name (**bond0**), **<network_interfaces>** represents the virtual functions (VFs) by their known name in the kernel and shown in the output of the **ip link** command (**eno1f0**, **eno2f0**), and *options* is a comma-separated list of bonding options. Enter **modinfo bonding** to see available options.
- When you create a bonded interface using **bond=**, you must specify how the IP address is assigned and other information for the bonded interface.
 - To configure the bonded interface to use DHCP, set the bond's IP address to **dhcp**. For example:

```
bond=bond0:eno1f0,eno2f0:mode=active-backup  
ip=bond0:dhcp
```

- To configure the bonded interface to use a static IP address, enter the specific IP address you want and related information. For example:

```
bond=bond0:eno1f0,eno2f0:mode=active-backup  
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:bond0:none
```

19.2.12.2. Installing RHCOS by using PXE booting

You can use PXE booting to install RHCOS on the machines.

Prerequisites

- You have created the Ignition config files for your cluster.
- You have configured suitable network, DNS and load balancing infrastructure.
- You have configured suitable PXE infrastructure.
- You have an HTTP server that can be accessed from your computer, and from the machines that you create.
- You have reviewed the *Advanced RHCOS installation configuration* section for different ways to configure features, such as networking and disk partitioning.

Procedure

1. Upload the bootstrap, control plane, and compute node Ignition config files that the installation program created to your HTTP server. Note the URLs of these files.



IMPORTANT

You can add or change configuration settings in your Ignition configs before saving them to your HTTP server. If you plan to add more compute machines to your cluster after you finish installation, do not delete these files.

- From the installation host, validate that the Ignition config files are available on the URLs. The following example gets the Ignition config file for the bootstrap node:

```
$ curl -k http://<HTTP_server>/bootstrap.ign 1
```

Example output

```
% Total    % Received % Xferd Average Speed   Time    Time     Time  Current
           Dload  Upload   Total   Spent    Left  Speed
  0   0   0   0   0   0   0   0   0  --:--:-- --:--:-- --:--:--    0{"ignition":
{"version":"3.2.0"},"passwd":{"users":[{"name":"core","sshAuthorizedKeys":["ssh-rsa...
```

Replace **bootstrap.ign** with **master.ign** or **worker.ign** in the command to validate that the Ignition config files for the control plane and compute nodes are also available.


- Although it is possible to obtain the RHCOS **kernel**, **initramfs** and **rootfs** files that are required for your preferred method of installing operating system instances from the [RHCOS image mirror](#) page, the recommended way to obtain the correct version of your RHCOS files are from the output of **openshift-install** command:

```
$ openshift-install coreos print-stream-json | grep -Eo "https.*(kernel-|initramfs-|rootfs-)\w+(\.img)?"
```

Example output

```
"<url>/art/storage/releases/rhcos-4.16-aarch64/<release>/aarch64/rhcos-<release>-live-
kernel-aarch64"
"<url>/art/storage/releases/rhcos-4.16-aarch64/<release>/aarch64/rhcos-<release>-live-
initramfs.aarch64.img"
"<url>/art/storage/releases/rhcos-4.16-aarch64/<release>/aarch64/rhcos-<release>-live-
rootfs.aarch64.img"
"<url>/art/storage/releases/rhcos-4.16-ppc64le/49.84.202110081256-0/ppc64le/rhcos-
<release>-live-kernel-ppc64le"
"<url>/art/storage/releases/rhcos-4.16-ppc64le/<release>/ppc64le/rhcos-<release>-live-
initramfs.ppc64le.img"
"<url>/art/storage/releases/rhcos-4.16-ppc64le/<release>/ppc64le/rhcos-<release>-live-
rootfs.ppc64le.img"
"<url>/art/storage/releases/rhcos-4.16-s390x/<release>/s390x/rhcos-<release>-live-kernel-
s390x"
"<url>/art/storage/releases/rhcos-4.16-s390x/<release>/s390x/rhcos-<release>-live-
initramfs.s390x.img"
"<url>/art/storage/releases/rhcos-4.16-s390x/<release>/s390x/rhcos-<release>-live-
rootfs.s390x.img"
"<url>/art/storage/releases/rhcos-4.16/<release>/x86_64/rhcos-<release>-live-kernel-
x86_64"
"<url>/art/storage/releases/rhcos-4.16/<release>/x86_64/rhcos-<release>-live-
```

```
initramfs.x86_64.img"
"<url>/art/storage/releases/rhcos-4.16/<release>/x86_64/rhcos-<release>-live-
rootfs.x86_64.img"
```




IMPORTANT

The RHCOS artifacts might not change with every release of OpenShift Container Platform. You must download images with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Only use the appropriate **kernel**, **initramfs**, and **rootfs** artifacts described below for this procedure. RHCOS QCOW2 images are not supported for this installation type.

The file names contain the OpenShift Container Platform version number. They resemble the following examples:

- **kernel:** rhcos-<version>-live-kernel-<architecture>
- **initramfs:** rhcos-<version>-live-initramfs.<architecture>.img
- **rootfs:** rhcos-<version>-live-rootfs.<architecture>.img

4. Upload the **rootfs**, **kernel**, and **initramfs** files to your HTTP server.



IMPORTANT

If you plan to add more compute machines to your cluster after you finish installation, do not delete these files.

5. Configure the network boot infrastructure so that the machines boot from their local disks after RHCOS is installed on them.
6. Configure PXE installation for the RHCOS images and begin the installation. Modify the following example menu entry for your environment and verify that the image and Ignition files are properly accessible:

```
DEFAULT pxeboot
TIMEOUT 20
PROMPT 0
LABEL pxeboot
  KERNEL http://<HTTP_server>/rhcos-<version>-live-kernel-<architecture> 1
  APPEND initrd=http://<HTTP_server>/rhcos-<version>-live-initramfs.<architecture>.img
  coreos.live.rootfs_url=http://<HTTP_server>/rhcos-<version>-live-rootfs.<architecture>.img
  coreos.inst.install_dev=/dev/sda coreos.inst.ignition_url=http://<HTTP_server>/bootstrap.ign
2 3
```

- 1** Specify the location of the live **kernel** file that you uploaded to your HTTP server. The URL must be HTTP, TFTP, or FTP; HTTPS and NFS are not supported.
- 2** If you use multiple NICs, specify a single interface in the **ip** option. For example, to use DHCP on a NIC that is named **eno1**, set **ip=eno1:dhcp**.

3


Specify the locations of the RHCOS files that you uploaded to your HTTP server. The **initrd** parameter value is the location of the **initramfs** file, the **coreos.live.rootfs_url**



NOTE

This configuration does not enable serial console access on machines with a graphical console. To configure a different console, add one or more **console=** arguments to the **APPEND** line. For example, add **console=tty0 console=ttyS0** to set the first PC serial port as the primary console and the graphical console as a secondary console. For more information, see [How does one set up a serial terminal and/or console in Red Hat Enterprise Linux?](#) and "Enabling the serial console for PXE and ISO installation" in the "Advanced RHCOS installation configuration" section.

7. Monitor the progress of the RHCOS installation on the console of the machine.



IMPORTANT


Be sure that the installation is successful on each node before commencing with the OpenShift Container Platform installation. Observing the installation process can also help to determine the cause of RHCOS installation issues that might arise.

8. After RHCOS installs, the system reboots. During reboot, the system applies the Ignition config file that you specified.
9. Check the console output to verify that Ignition ran.

Example command

```
Ignition: ran on 2022/03/14 14:48:33 UTC (this boot)
Ignition: user-provided config was applied
```

10. Continue to create the machines for your cluster.



IMPORTANT

You must create the bootstrap and control plane machines at this time. If the control plane machines are not made schedulable, also create at least two compute machines before you install the cluster.

If the required network, DNS, and load balancer infrastructure are in place, the OpenShift Container Platform bootstrap process begins automatically after the RHCOS nodes have rebooted.



NOTE

RHCOS nodes do not include a default password for the **core** user. You can access the nodes by running **ssh core@<node>.<cluster_name>.<base_domain>** as a user with access to the SSH private key that is paired to the public key that you specified in your **install_config.yaml** file. OpenShift Container Platform 4 cluster nodes running RHCOS are immutable and rely on Operators to apply cluster changes. Accessing cluster nodes by using SSH is not recommended. However, when investigating installation issues, if the OpenShift Container Platform API is not available, or the kubelet is not properly functioning on a target node, SSH access might be required for debugging or disaster recovery.

19.2.12.3. Enabling multipathing with kernel arguments on RHCOS

In OpenShift Container Platform version 4.16, during installation, you can enable multipathing for provisioned nodes. RHCOS supports multipathing on the primary disk. Multipathing provides added benefits of stronger resilience to hardware failure to achieve higher host availability.

During the initial cluster creation, you might want to add kernel arguments to all master or worker nodes. To add kernel arguments to master or worker nodes, you can create a **MachineConfig** object and inject that object into the set of manifest files used by Ignition during cluster setup.

Procedure

1. Change to the directory that contains the installation program and generate the Kubernetes manifests for the cluster:

```
$ ./openshift-install create manifests --dir <installation_directory>
```

2. Decide if you want to add kernel arguments to worker or control plane nodes.
 - Create a machine config file. For example, create a **99-master-kargs-mpath.yaml** that instructs the cluster to add the **master** label and identify the multipath kernel argument:

```
apiVersion: machineconfiguration.openshift.io/v1
kind: MachineConfig
metadata:
  labels:
    machineconfiguration.openshift.io/role: "master"
  name: 99-master-kargs-mpath
spec:
  kernelArguments:
    - 'rd.multipath=default'
    - 'root=/dev/disk/by-label/dm-mpath-root'
```

3. To enable multipathing on worker nodes:
 - Create a machine config file. For example, create a **99-worker-kargs-mpath.yaml** that instructs the cluster to add the **worker** label and identify the multipath kernel argument:

```
apiVersion: machineconfiguration.openshift.io/v1
kind: MachineConfig
metadata:
  labels:
```



```

machineconfiguration.openshift.io/role: "worker"
name: 99-worker-kargs-mpath
spec:
  kernelArguments:
    - 'rd.multipath=default'
    - 'root=/dev/disk/by-label/dm-mpath-root'

```

You can now continue on to create the cluster.



IMPORTANT

Additional postinstallation steps are required to fully enable multipathing. For more information, see “Enabling multipathing with kernel arguments on RHCOS” in *Postinstallation machine configuration tasks*.

In case of MPIO failure, use the `bootlist` command to update the boot device list with alternate logical device names. The command displays a boot list and it designates the possible boot devices for when the system is booted in normal mode.

- a. To display a boot list and specify the possible boot devices if the system is booted in normal mode, enter the following command:

```

$ bootlist -m normal -o
sda

```

- b. To update the boot list for normal mode and add alternate device names, enter the following command:

```

$ bootlist -m normal -o /dev/sdc /dev/sdd /dev/sde
sdc
sdd
sde

```

If the original boot disk path is down, the node reboots from the alternate device registered in the normal boot device list.

19.2.13. Waiting for the bootstrap process to complete

The OpenShift Container Platform bootstrap process begins after the cluster nodes first boot into the persistent RHCOS environment that has been installed to disk. The configuration information provided through the Ignition config files is used to initialize the bootstrap process and install OpenShift Container Platform on the machines. You must wait for the bootstrap process to complete.

Prerequisites

- You have created the Ignition config files for your cluster.
- You have configured suitable network, DNS and load balancing infrastructure.
- You have obtained the installation program and generated the Ignition config files for your cluster.
- You installed RHCOS on your cluster machines and provided the Ignition config files that the OpenShift Container Platform installation program generated.

- Your machines have direct internet access or have an HTTP or HTTPS proxy available.

Procedure

- Monitor the bootstrap process:

```
$ ./openshift-install --dir <installation_directory> wait-for bootstrap-complete \ 1
--log-level=info 2
```

- For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

- To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Example output

```
INFO Waiting up to 30m0s for the Kubernetes API at https://api.test.example.com:6443...
INFO API v1.29.4 up
INFO Waiting up to 30m0s for bootstrapping to complete...
INFO It is now safe to remove the bootstrap resources
```

The command succeeds when the Kubernetes API server signals that it has been bootstrapped on the control plane machines.

- After the bootstrap process is complete, remove the bootstrap machine from the load balancer.



IMPORTANT

You must remove the bootstrap machine from the load balancer at this point. You can also remove or reformat the bootstrap machine itself.

19.2.14. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

- Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

19.2.15. Approving the certificate signing requests for your machines

When you add machines to a cluster, two pending certificate signing requests (CSRs) are generated for each machine that you added. You must confirm that these CSRs are approved or, if necessary, approve them yourself. The client requests must be approved first, followed by the server requests.

Prerequisites

- You added machines to your cluster.

Procedure

1. Confirm that the cluster recognizes the machines:

```
$ oc get nodes
```

Example output

```
NAME      STATUS   ROLES    AGE   VERSION
master-0  Ready   master   63m   v1.29.4
master-1  Ready   master   63m   v1.29.4
master-2  Ready   master   64m   v1.29.4
```

The output lists all of the machines that you created.



NOTE

The preceding output might not include the compute nodes, also known as worker nodes, until some CSRs are approved.

2. Review the pending CSRs and ensure that you see the client requests with the **Pending** or **Approved** status for each machine that you added to the cluster:


```
$ oc get csr
```

Example output

```
NAME      AGE   REQUESTOR                                     CONDITION
csr-8b2br  15m   system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper  Pending
csr-8vnps  15m   system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper  Pending
...
```

In this example, two machines are joining the cluster. You might see more approved CSRs in the list.

3. If the CSRs were not approved, after all of the pending CSRs for the machines you added are in **Pending** status, approve the CSRs for your cluster machines:



NOTE

Because the CSRs rotate automatically, approve your CSRs within an hour of adding the machines to the cluster. If you do not approve them within an hour, the certificates will rotate, and more than two certificates will be present for each node. You must approve all of these certificates. After the client CSR is approved, the Kubelet creates a secondary CSR for the serving certificate, which requires manual approval. Then, subsequent serving certificate renewal requests are automatically approved by the **machine-approver** if the Kubelet requests a new certificate with identical parameters.



NOTE

For clusters running on platforms that are not machine API enabled, such as bare metal and other user-provisioned infrastructure, you must implement a method of automatically approving the kubelet serving certificate requests (CSRs). If a request is not approved, then the **oc exec**, **oc rsh**, and **oc logs** commands cannot succeed, because a serving certificate is required when the API server connects to the kubelet. Any operation that contacts the Kubelet endpoint requires this certificate approval to be in place. The method must watch for new CSRs, confirm that the CSR was submitted by the **node-bootstrap** service account in the **system:node** or **system:admin** groups, and confirm the identity of the node.

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{"\n"}\n{{end}}{{end}}' | xargs --no-run-if-empty oc adm certificate approve
```



NOTE

Some Operators might not become available until some CSRs are approved.

4. Now that your client requests are approved, you must review the server requests for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

```

NAME      AGE    REQUESTOR                                CONDITION
csr-bfd72 5m26s system:node:ip-10-0-50-126.us-east-2.compute.internal
Pending
csr-c57lv 5m26s system:node:ip-10-0-95-157.us-east-2.compute.internal
Pending
...

```

5. If the remaining CSRs are not approved, and are in the **Pending** status, approve the CSRs for your cluster machines:

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

1 **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{\n"}\n{{end}}\n{{end}}' | xargs oc adm certificate approve
```

6. After all client and server CSRs have been approved, the machines have the **Ready** status. Verify this by running the following command:

```
$ oc get nodes
```

Example output

```

NAME      STATUS  ROLES  AGE  VERSION
master-0  Ready   master 73m  v1.29.4
master-1  Ready   master 73m  v1.29.4
master-2  Ready   master 74m  v1.29.4
worker-0  Ready   worker 11m  v1.29.4
worker-1  Ready   worker 11m  v1.29.4

```



NOTE

It can take a few minutes after approval of the server CSRs for the machines to transition to the **Ready** status.

Additional information

- For more information on CSRs, see [Certificate Signing Requests](#).

19.2.16. Initial Operator configuration

After the control plane initializes, you must immediately configure some Operators so that they all become available.

Prerequisites

- Your control plane has initialized.

Procedure

1. Watch the cluster components come online:

```
$ watch -n5 oc get clusteroperators
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.16.0	True	False	False	19m
baremetal	4.16.0	True	False	False	37m
cloud-credential	4.16.0	True	False	False	40m
cluster-autoscaler	4.16.0	True	False	False	37m
config-operator	4.16.0	True	False	False	38m
console	4.16.0	True	False	False	26m
csi-snapshot-controller	4.16.0	True	False	False	37m
dns	4.16.0	True	False	False	37m
etcd	4.16.0	True	False	False	36m
image-registry	4.16.0	True	False	False	31m
ingress	4.16.0	True	False	False	30m
insights	4.16.0	True	False	False	31m
kube-apiserver	4.16.0	True	False	False	26m
kube-controller-manager	4.16.0	True	False	False	36m
kube-scheduler	4.16.0	True	False	False	36m
kube-storage-version-migrator	4.16.0	True	False	False	37m
machine-api	4.16.0	True	False	False	29m
machine-approver	4.16.0	True	False	False	37m
machine-config	4.16.0	True	False	False	36m
marketplace	4.16.0	True	False	False	37m
monitoring	4.16.0	True	False	False	29m
network	4.16.0	True	False	False	38m
node-tuning	4.16.0	True	False	False	37m
openshift-apiserver	4.16.0	True	False	False	32m
openshift-controller-manager	4.16.0	True	False	False	30m
openshift-samples	4.16.0	True	False	False	32m
operator-lifecycle-manager	4.16.0	True	False	False	37m
operator-lifecycle-manager-catalog	4.16.0	True	False	False	37m
operator-lifecycle-manager-packageserver	4.16.0	True	False	False	32m
service-ca	4.16.0	True	False	False	38m
storage	4.16.0	True	False	False	37m

2. Configure the Operators that are not available.

19.2.16.1. Image registry storage configuration

The Image Registry Operator is not initially available for platforms that do not provide default storage. After installation, you must configure your registry to use storage so that the Registry Operator is made available.

Instructions are shown for configuring a persistent volume, which is required for production clusters. Where applicable, instructions are shown for configuring an empty directory as the storage location, which is available for only non-production clusters.


Additional instructions are provided for allowing the image registry to use block storage types by using the **Recreate** rollout strategy during upgrades.

19.2.16.1.1. Configuring registry storage for IBM Power

As a cluster administrator, following installation you must configure your registry to use storage.

Prerequisites

- You have access to the cluster as a user with the **cluster-admin** role.
- You have a cluster on IBM Power®.
- You have provisioned persistent storage for your cluster, such as Red Hat OpenShift Data Foundation.



IMPORTANT

OpenShift Container Platform supports **ReadWriteOnce** access for image registry storage when you have only one replica. **ReadWriteOnce** access also requires that the registry uses the **Recreate** rollout strategy. To deploy an image registry that supports high availability with two or more replicas, **ReadWriteMany** access is required.

- Must have 100Gi capacity.

Procedure

1. To configure your registry to use storage, change the **spec.storage.pvc** in the **configs.imageregistry/cluster** resource.



NOTE

When you use shared storage, review your security settings to prevent outside access.

2. Verify that you do not have a registry pod:

```
$ oc get pod -n openshift-image-registry -l docker-registry=default
```

Example output

```
No resources found in openshift-image-registry namespace
```



NOTE

If you do have a registry pod in your output, you do not need to continue with this procedure.

3. Check the registry configuration:

```
$ oc edit configs.imageregistry.operator.openshift.io
```

Example output

```
storage:
  pvc:
    claim:
```

Leave the **claim** field blank to allow the automatic creation of an **image-registry-storage** PVC.

4. Check the **clusteroperator** status:

```
$ oc get clusteroperator image-registry
```

Example output

```
NAME          VERSION          AVAILABLE PROGRESSING DEGRADED SINCE
MESSAGE
image-registry 4.16             True      False      False  6h50m
```

5. Ensure that your registry is set to managed to enable building and pushing of images.

- Run:

```
$ oc edit configs.imageregistry/cluster
```

Then, change the line

```
managementState: Removed
```

to

```
managementState: Managed
```

19.2.16.1.2. Configuring storage for the image registry in non-production clusters

You must configure storage for the Image Registry Operator. For non-production clusters, you can set the image registry to an empty directory. If you do so, all images are lost if you restart the registry.

Procedure

- To set the image registry storage to an empty directory:

```
$ oc patch configs.imageregistry.operator.openshift.io cluster --type merge --patch '{"spec": {"storage":{"emptyDir":{}}}'
```



WARNING

Configure this option for only non-production clusters.

If you run this command before the Image Registry Operator initializes its components, the **oc patch** command fails with the following error:

```
Error from server (NotFound): configs.imageregistry.operator.openshift.io "cluster" not found
```

Wait a few minutes and run the command again.

19.2.17. Completing installation on user-provisioned infrastructure

After you complete the Operator configuration, you can finish installing the cluster on infrastructure that you provide.

Prerequisites

- Your control plane has initialized.
- You have completed the initial Operator configuration.

Procedure

1. Confirm that all the cluster components are online with the following command:

```
$ watch -n5 oc get clusteroperators
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.16.0	True	False	False	19m
baremetal	4.16.0	True	False	False	37m
cloud-credential	4.16.0	True	False	False	40m
cluster-autoscaler	4.16.0	True	False	False	37m
config-operator	4.16.0	True	False	False	38m
console	4.16.0	True	False	False	26m
csi-snapshot-controller	4.16.0	True	False	False	37m
dns	4.16.0	True	False	False	37m
etcd	4.16.0	True	False	False	36m
image-registry	4.16.0	True	False	False	31m
ingress	4.16.0	True	False	False	30m
insights	4.16.0	True	False	False	31m
kube-apiserver	4.16.0	True	False	False	26m
kube-controller-manager	4.16.0	True	False	False	36m
kube-scheduler	4.16.0	True	False	False	36m
kube-storage-version-migrator	4.16.0	True	False	False	37m
machine-api	4.16.0	True	False	False	29m
machine-approver	4.16.0	True	False	False	37m
machine-config	4.16.0	True	False	False	36m
marketplace	4.16.0	True	False	False	37m
monitoring	4.16.0	True	False	False	29m
network	4.16.0	True	False	False	38m
node-tuning	4.16.0	True	False	False	37m
openshift-apiserver	4.16.0	True	False	False	32m
openshift-controller-manager	4.16.0	True	False	False	30m
openshift-samples	4.16.0	True	False	False	32m

```
operator-lifecycle-manager      4.16.0 True    False   False   37m
operator-lifecycle-manager-catalog  4.16.0 True    False   False   37m
operator-lifecycle-manager-packageserver 4.16.0 True    False   False   32m
service-ca                      4.16.0 True    False   False   38m
storage                          4.16.0 True    False   False   37m
```

Alternatively, the following command notifies you when all of the clusters are available. It also retrieves and displays credentials:

```
$ ./openshift-install --dir <installation_directory> wait-for install-complete 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

Example output

```
INFO Waiting up to 30m0s for the cluster to initialize...
```

The command succeeds when the Cluster Version Operator finishes deploying the OpenShift Container Platform cluster from Kubernetes API server.

IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

2. Confirm that the Kubernetes API server is communicating with the pods.

- a. To view a list of all pods, use the following command:

```
$ oc get pods --all-namespaces
```

Example output

```
NAMESPACE          NAME                                     READY STATUS
RESTARTS  AGE
openshift-apiserver-operator  openshift-apiserver-operator-85cb746d55-zqhs8  1/1
Running    1    9m
openshift-apiserver          apiserver-67b9g                                1/1  Running  0
3m
openshift-apiserver          apiserver-ljcmx                                1/1  Running  0
```

```

1m
openshift-apiserver          apiserver-z25h4           1/1   Running   0
2m
openshift-authentication-operator authentication-operator-69d5d8bf84-vh2n8 1/1
Running   0       5m
...

```

- b. View the logs for a pod that is listed in the output of the previous command by using the following command:

```
$ oc logs <pod_name> -n <namespace> ❶
```

- ❶ Specify the pod name and namespace, as shown in the output of the previous command.

If the pod logs display, the Kubernetes API server can communicate with the cluster machines.

3. Additional steps are required to enable multipathing. Do not enable multipathing during installation.
See "Enabling multipathing with kernel arguments on RHCOS" in the *Postinstallation machine configuration tasks* documentation for more information.

19.2.18. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service

19.2.19. Next steps

- [Enabling multipathing with kernel arguments on RHCOS](#) .
- [Customize your cluster](#).
- If necessary, you can [opt out of remote health reporting](#) .

19.3. INSTALLING A CLUSTER ON IBM POWER IN A RESTRICTED NETWORK

In OpenShift Container Platform version 4.16, you can install a cluster on IBM Power® infrastructure that you provision in a restricted network.



IMPORTANT

Additional considerations exist for non-bare metal platforms. Review the information in the [guidelines for deploying OpenShift Container Platform on non-tested platforms](#) before you install an OpenShift Container Platform cluster.

19.3.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [created a mirror registry for installation in a restricted network](#) and obtained the **imageContentSources** data for your version of OpenShift Container Platform.
- Before you begin the installation process, you must move or remove any existing installation files. This ensures that the required installation files are created and updated during the installation process.



IMPORTANT

Ensure that installation steps are performed on a machine with access to the installation media.

- You provisioned [persistent storage using OpenShift Data Foundation](#) or other supported storage protocols for your cluster. To deploy a private image registry, you must set up persistent storage with **ReadWriteMany** access.
- If you use a firewall and plan to use the Telemetry service, you [configured the firewall to allow the sites](#) that your cluster requires access to.



NOTE

Be sure to also review this site list if you are configuring a proxy.

19.3.2. About installations in restricted networks

In OpenShift Container Platform 4.16, you can perform an installation that does not require an active connection to the internet to obtain software components. Restricted network installations can be completed using installer-provisioned infrastructure or user-provisioned infrastructure, depending on the cloud platform to which you are installing the cluster.

To complete a restricted network installation, you must create a registry that mirrors the contents of the OpenShift image registry and contains the installation media. You can create this registry on a mirror host, which can access both the internet and your closed network, or by using other methods that meet your restrictions.



IMPORTANT

Because of the complexity of the configuration for user-provisioned installations, consider completing a standard user-provisioned infrastructure installation before you attempt a restricted network installation using user-provisioned infrastructure. Completing this test installation might make it easier to isolate and troubleshoot any issues that might arise during your installation in a restricted network.

19.3.2.1. Additional limits

Clusters in restricted networks have the following additional limitations and restrictions:

- The **ClusterVersion** status includes an **Unable to retrieve available updates** error.
- By default, you cannot use the contents of the Developer Catalog because you cannot access the required image stream tags.

19.3.3. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to obtain the images that are necessary to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.

19.3.4. Requirements for a cluster with user-provisioned infrastructure

For a cluster that contains user-provisioned infrastructure, you must deploy all of the required machines.

This section describes the requirements for deploying OpenShift Container Platform on user-provisioned infrastructure.

19.3.4.1. Required machines for cluster installation

The smallest OpenShift Container Platform clusters require the following hosts:

Table 19.20. Minimum required hosts

Hosts	Description
One temporary bootstrap machine	The cluster requires the bootstrap machine to deploy the OpenShift Container Platform cluster on the three control plane machines. You can remove the bootstrap machine after you install the cluster.

Hosts	Description
Three control plane machines	The control plane machines run the Kubernetes and OpenShift Container Platform services that form the control plane.
At least two compute machines, which are also known as worker machines.	The workloads requested by OpenShift Container Platform users run on the compute machines.



IMPORTANT

To maintain high availability of your cluster, use separate physical hosts for these cluster machines.

The bootstrap, control plane, and compute machines must use Red Hat Enterprise Linux CoreOS (RHCOS) as the operating system.

Note that RHCOS is based on Red Hat Enterprise Linux (RHEL) 9.2 and inherits all of its hardware certifications and requirements. See [Red Hat Enterprise Linux technology capabilities and limits](#) .

19.3.4.2. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 19.21. Minimum resource requirements

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Bootstrap	RHCOS	2	16 GB	100 GB	300
Control plane	RHCOS	2	16 GB	100 GB	300
Compute	RHCOS	2	8 GB	100 GB	300

1. One vCPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio: $(\text{threads per core} \times \text{cores}) \times \text{sockets} = \text{vCPUs}$.
2. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes. Note that on many cloud platforms, storage size and IOPS scale together, so you might need to over-allocate storage volume to obtain sufficient performance.

**NOTE**

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

19.3.4.3. Minimum IBM Power requirements

You can install OpenShift Container Platform version 4.16 on the following IBM® hardware:

- IBM Power®9 or IBM Power®10 processor-based systems

**NOTE**

Support for RHCOS functionality for all IBM Power®8 models, IBM Power® AC922, IBM Power® IC922, and IBM Power® LC922 is deprecated in OpenShift Container Platform 4.16. Red Hat recommends that you use later hardware models.

Hardware requirements

- Six logical partitions (LPARs) across multiple PowerVM servers

Operating system requirements

- One instance of an IBM Power®9 or Power10 processor-based system

On your IBM Power® instance, set up:

- Three LPARs for OpenShift Container Platform control plane machines
- Two LPARs for OpenShift Container Platform compute machines
- One LPAR for the temporary OpenShift Container Platform bootstrap machine

Disk storage for the IBM Power guest virtual machines

- Local storage, or storage provisioned by the Virtual I/O Server using vSCSI, NPIV (N-Port ID Virtualization) or SSP (shared storage pools)

Network for the PowerVM guest virtual machines

- Dedicated physical adapter, or SR-IOV virtual function
- Available by the Virtual I/O Server using Shared Ethernet Adapter
- Virtualized by the Virtual I/O Server using IBM® vNIC

Storage / main memory

- 100 GB / 16 GB for OpenShift Container Platform control plane machines
- 100 GB / 8 GB for OpenShift Container Platform compute machines
- 100 GB / 16 GB for the temporary OpenShift Container Platform bootstrap machine

19.3.4.4. Recommended IBM Power system requirements**Hardware requirements**

- Six LPARs across multiple PowerVM servers

Operating system requirements

- One instance of an IBM Power®9 or IBM Power®10 processor-based system

On your IBM Power® instance, set up:

- Three LPARs for OpenShift Container Platform control plane machines
- Two LPARs for OpenShift Container Platform compute machines
- One LPAR for the temporary OpenShift Container Platform bootstrap machine

Disk storage for the IBM Power guest virtual machines

- Local storage, or storage provisioned by the Virtual I/O Server using vSCSI, NPIV (N-Port ID Virtualization) or SSP (shared storage pools)

Network for the PowerVM guest virtual machines

- Dedicated physical adapter, or SR-IOV virtual function
- Virtualized by the Virtual I/O Server using Shared Ethernet Adapter
- Virtualized by the Virtual I/O Server using IBM® vNIC

Storage / main memory

- 120 GB / 32 GB for OpenShift Container Platform control plane machines
- 120 GB / 32 GB for OpenShift Container Platform compute machines
- 120 GB / 16 GB for the temporary OpenShift Container Platform bootstrap machine

19.3.4.5. Certificate signing requests management

Because your cluster has limited access to automatic machine management when you use infrastructure

that you provision, you must provide a mechanism for approving cluster certificate signing requests (CSRs) after installation. The **kube-controller-manager** only approves the kubelet client CSRs. The **machine-approver** cannot guarantee the validity of a serving certificate that is requested by using kubelet credentials because it cannot confirm that the correct machine issued the request. You must determine and implement a method of verifying the validity of the kubelet serving certificate requests and approving them.

19.3.4.6. Networking requirements for user-provisioned infrastructure

All the Red Hat Enterprise Linux CoreOS (RHCOS) machines require networking to be configured in **initramfs** during boot to fetch their Ignition config files.

During the initial boot, the machines require an IP address configuration that is set either through a DHCP server or statically by providing the required boot options. After a network connection is established, the machines download their Ignition config files from an HTTP or HTTPS server. The Ignition config files are then used to set the exact state of each machine. The Machine Config Operator completes more changes to the machines, such as the application of new certificates or keys, after installation.

It is recommended to use a DHCP server for long-term management of the cluster machines. Ensure that the DHCP server is configured to provide persistent IP addresses, DNS server information, and hostnames to the cluster machines.



NOTE

If a DHCP service is not available for your user-provisioned infrastructure, you can instead provide the IP networking configuration and the address of the DNS server to the nodes at RHCOS install time. These can be passed as boot arguments if you are installing from an ISO image. See the *Installing RHCOS and starting the OpenShift Container Platform bootstrap process* section for more information about static IP provisioning and advanced networking options.

The Kubernetes API server must be able to resolve the node names of the cluster machines. If the API servers and worker nodes are in different zones, you can configure a default DNS search zone to allow the API server to resolve the node names. Another supported approach is to always refer to hosts by their fully-qualified domain names in both the node objects and all DNS requests.

19.3.4.6.1. Setting the cluster node hostnames through DHCP

On Red Hat Enterprise Linux CoreOS (RHCOS) machines, the hostname is set through NetworkManager. By default, the machines obtain their hostname through DHCP. If the hostname is not provided by DHCP, set statically through kernel arguments, or another method, it is obtained through a reverse DNS lookup. Reverse DNS lookup occurs after the network has been initialized on a node and can take time to resolve. Other system services can start prior to this and detect the hostname as **localhost** or similar. You can avoid this by using DHCP to provide the hostname for each cluster node.

Additionally, setting the hostnames through DHCP can bypass any manual DNS record name configuration errors in environments that have a DNS split-horizon implementation.

19.3.4.6.2. Network connectivity requirements

You must configure the network connectivity between machines to allow OpenShift Container Platform cluster components to communicate. Each machine must be able to resolve the hostnames of all other machines in the cluster.

This section provides details about the ports that are required.

Table 19.22. Ports used for all-machine to all-machine communications

Protocol	Port	Description
ICMP	N/A	Network reachability tests
TCP	1936	Metrics
	9000-9999	Host level services, including the node exporter on ports 9100-9101 and the Cluster Version Operator on port 9099 .
	10250-10259	The default ports that Kubernetes reserves
UDP	4789	VXLAN
	6081	Geneve
	9000-9999	Host level services, including the node exporter on ports 9100-9101 .
	500	IPsec IKE packets
	4500	IPsec NAT-T packets
	123	Network Time Protocol (NTP) on UDP port 123 If an external NTP time server is configured, you must open UDP port 123 .
TCP/UDP	30000-32767	Kubernetes node port
ESP	N/A	IPsec Encapsulating Security Payload (ESP)

Table 19.23. Ports used for all-machine to control plane communications

Protocol	Port	Description
TCP	6443	Kubernetes API

Table 19.24. Ports used for control plane machine to control plane machine communications

Protocol	Port	Description
TCP	2379-2380	etcd server and peer ports

NTP configuration for user-provisioned infrastructure

OpenShift Container Platform clusters are configured to use a public Network Time Protocol (NTP) server by default. If you want to use a local enterprise NTP server, or if your cluster is being deployed in a disconnected network, you can configure the cluster to use a specific time server. For more information, see the documentation for *Configuring chrony time service*.

If a DHCP server provides NTP server information, the chrony time service on the Red Hat Enterprise Linux CoreOS (RHCOS) machines read the information and can sync the clock with the NTP servers.

Additional resources

- [Configuring chrony time service](#)


19.3.4.7. User-provisioned DNS requirements

In OpenShift Container Platform deployments, DNS name resolution is required for the following components:

- The Kubernetes API
- The OpenShift Container Platform application wildcard
- The bootstrap, control plane, and compute machines

Reverse DNS resolution is also required for the Kubernetes API, the bootstrap machine, the control plane machines, and the compute machines.

DNS A/AAAA or CNAME records are used for name resolution and PTR records are used for reverse name resolution. The reverse records are important because Red Hat Enterprise Linux CoreOS (RHCOS) uses the reverse records to set the hostnames for all the nodes, unless the hostnames are provided by DHCP. Additionally, the reverse records are used to generate the certificate signing requests (CSR) that OpenShift Container Platform needs to operate.




NOTE

It is recommended to use a DHCP server to provide the hostnames to each cluster node. See the *DHCP recommendations for user-provisioned infrastructure* section for more information.

The following DNS records are required for a user-provisioned OpenShift Container Platform cluster and they must be in place before installation. In each record, **<cluster_name>** is the cluster name and **<base_domain>** is the base domain that you specify in the **install-config.yaml** file. A complete DNS record takes the form: **<component>.<cluster_name>.<base_domain>..**

Table 19.25. Required DNS records

Component	Record	Description
Kubernetes API	api.<cluster_name>.<base_domain>..	A DNS A/AAAA or CNAME record, and a DNS PTR record, to identify the API load balancer. These records must be resolvable by both clients external to the cluster and from all the nodes within the cluster.

Component	Record	Description
	api-int.<cluster_name>.<base_domain>.	<p>A DNS A/AAAA or CNAME record, and a DNS PTR record, to internally identify the API load balancer. These records must be resolvable from all the nodes within the cluster.</p> <div style="display: flex; align-items: flex-start;">  <div> <p>IMPORTANT</p> <p>The API server must be able to resolve the worker nodes by the hostnames that are recorded in Kubernetes. If the API server cannot resolve the node names, then proxied API calls can fail, and you cannot retrieve logs from pods.</p> </div> </div>
Routes	*.apps.<cluster_name>.<base_domain>.	<p>A wildcard DNS A/AAAA or CNAME record that refers to the application ingress load balancer. The application ingress load balancer targets the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default. These records must be resolvable by both clients external to the cluster and from all the nodes within the cluster.</p> <p>For example, console-openshift-console.apps.<cluster_name>.<base_domain> is used as a wildcard route to the OpenShift Container Platform console.</p>
Bootstrap machine	bootstrap.<cluster_name>.<base_domain>.	A DNS A/AAAA or CNAME record, and a DNS PTR record, to identify the bootstrap machine. These records must be resolvable by the nodes within the cluster.
Control plane machines	<control_plane><n>.<cluster_name>.<base_domain>.	DNS A/AAAA or CNAME records and DNS PTR records to identify each machine for the control plane nodes. These records must be resolvable by the nodes within the cluster.
Compute machines	<compute><n>.<cluster_name>.<base_domain>.	DNS A/AAAA or CNAME records and DNS PTR records to identify each machine for the worker nodes. These records must be resolvable by the nodes within the cluster.



NOTE

In OpenShift Container Platform 4.4 and later, you do not need to specify etcd host and SRV records in your DNS configuration.

TIP

You can use the **dig** command to verify name and reverse name resolution. See the section on *Validating DNS resolution for user-provisioned infrastructure* for detailed validation steps.

19.3.4.7.1. Example DNS configuration for user-provisioned clusters

This section provides A and PTR record configuration samples that meet the DNS requirements for deploying OpenShift Container Platform on user-provisioned infrastructure. The samples are not meant to provide advice for choosing one DNS solution over another.

In the examples, the cluster name is **ocp4** and the base domain is **example.com**.

Example DNS A record configuration for a user-provisioned cluster

The following example is a BIND zone file that shows sample A records for name resolution in a user-provisioned cluster.

Example 19.4. Sample DNS zone database

```
$TTL 1W
@ IN SOA ns1.example.com. root (
    2019070700 ; serial
    3H ; refresh (3 hours)
    30M ; retry (30 minutes)
    2W ; expiry (2 weeks)
    1W ) ; minimum (1 week)
IN NS ns1.example.com.
IN MX 10 smtp.example.com.
;
;
ns1.example.com. IN A 192.168.1.5
smtp.example.com. IN A 192.168.1.5
;
helper.example.com. IN A 192.168.1.5
helper.ocp4.example.com. IN A 192.168.1.5
;
api.ocp4.example.com. IN A 192.168.1.5 ①
api-int.ocp4.example.com. IN A 192.168.1.5 ②
;
*.apps.ocp4.example.com. IN A 192.168.1.5 ③
;
bootstrap.ocp4.example.com. IN A 192.168.1.96 ④
;
control-plane0.ocp4.example.com. IN A 192.168.1.97 ⑤
control-plane1.ocp4.example.com. IN A 192.168.1.98 ⑥
control-plane2.ocp4.example.com. IN A 192.168.1.99 ⑦
;
compute0.ocp4.example.com. IN A 192.168.1.11 ⑧
compute1.ocp4.example.com. IN A 192.168.1.7 ⑨
;
;EOF
```

- ① Provides name resolution for the Kubernetes API. The record refers to the IP address of the API load balancer.

- 2 Provides name resolution for the Kubernetes API. The record refers to the IP address of the API load balancer and is used for internal cluster communications.
- 3 Provides name resolution for the wildcard routes. The record refers to the IP address of the application ingress load balancer. The application ingress load balancer targets the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.



NOTE

In the example, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

- 4 Provides name resolution for the bootstrap machine.
- 5 6 7 Provides name resolution for the control plane machines.
- 8 9 Provides name resolution for the compute machines.

Example DNS PTR record configuration for a user-provisioned cluster

The following example BIND zone file shows sample PTR records for reverse name resolution in a user-provisioned cluster.

Example 19.5. Sample DNS zone database for reverse records

```
$TTL 1W
@ IN SOA ns1.example.com. root (
    2019070700 ; serial
    3H ; refresh (3 hours)
    30M ; retry (30 minutes)
    2W ; expiry (2 weeks)
    1W ) ; minimum (1 week)
IN NS ns1.example.com.
;
5.1.168.192.in-addr.arpa. IN PTR api.ocp4.example.com. 1
5.1.168.192.in-addr.arpa. IN PTR api-int.ocp4.example.com. 2
;
96.1.168.192.in-addr.arpa. IN PTR bootstrap.ocp4.example.com. 3
;
97.1.168.192.in-addr.arpa. IN PTR control-plane0.ocp4.example.com. 4
98.1.168.192.in-addr.arpa. IN PTR control-plane1.ocp4.example.com. 5
99.1.168.192.in-addr.arpa. IN PTR control-plane2.ocp4.example.com. 6
;
11.1.168.192.in-addr.arpa. IN PTR compute0.ocp4.example.com. 7
7.1.168.192.in-addr.arpa. IN PTR compute1.ocp4.example.com. 8
;
;EOF
```

- 1 Provides reverse DNS resolution for the Kubernetes API. The PTR record refers to the record name of the API load balancer.
- 2 Provides reverse DNS resolution for the Kubernetes API. The PTR record refers to the record name of the API load balancer and is used for internal cluster communications.
- 3 Provides reverse DNS resolution for the bootstrap machine.
- 4 5 6 Provides reverse DNS resolution for the control plane machines.
- 7 8 Provides reverse DNS resolution for the compute machines.

**NOTE**

A PTR record is not required for the OpenShift Container Platform application wildcard.

19.3.4.8. Load balancing requirements for user-provisioned infrastructure


Before you install OpenShift Container Platform, you must provision the API and application Ingress load balancing infrastructure. In production scenarios, you can deploy the API and application Ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

**NOTE**

If you want to deploy the API and application Ingress load balancers with a Red Hat Enterprise Linux (RHEL) instance, you must purchase the RHEL subscription separately.

The load balancing infrastructure must meet the following requirements:

1. **API load balancer.** Provides a common endpoint for users, both human and machine, to interact with and configure the platform. Configure the following conditions:
 - Layer 4 load balancing only. This can be referred to as Raw TCP or SSL Passthrough mode.
 - A stateless load balancing algorithm. The options vary based on the load balancer implementation.

**IMPORTANT**

Do not configure session persistence for an API load balancer. Configuring session persistence for a Kubernetes API server might cause performance issues from excess application traffic for your OpenShift Container Platform cluster and the Kubernetes API that runs inside the cluster.

Configure the following ports on both the front and back of the load balancers:

Table 19.26. API load balancer

Port	Back-end machines (pool members)	Internal	External	Description
------	----------------------------------	----------	----------	-------------

Port	Back-end machines (pool members)	Internal	External	Description
6443	Bootstrap and control plane. You remove the bootstrap machine from the load balancer after the bootstrap machine initializes the cluster control plane. You must configure the /readyz endpoint for the API server health check probe.	X	X	Kubernetes API server
22623	Bootstrap and control plane. You remove the bootstrap machine from the load balancer after the bootstrap machine initializes the cluster control plane.	X		Machine config server



NOTE

The load balancer must be configured to take a maximum of 30 seconds from the time the API server turns off the **/readyz** endpoint to the removal of the API server instance from the pool. Within the time frame after **/readyz** returns an error or becomes healthy, the endpoint must have been removed or added. Probing every 5 or 10 seconds, with two successful requests to become healthy and three to become unhealthy, are well-tested values.

2. **Application Ingress load balancer:** Provides an ingress point for application traffic flowing in from outside the cluster. A working configuration for the Ingress router is required for an OpenShift Container Platform cluster.

Configure the following conditions:

- Layer 4 load balancing only. This can be referred to as Raw TCP or SSL Passthrough mode.
- A connection-based or session-based persistence is recommended, based on the options available and types of applications that will be hosted on the platform.

TIP


If the true IP address of the client can be seen by the application Ingress load balancer, enabling source IP-based session persistence can improve performance for applications that use end-to-end TLS encryption.

Configure the following ports on both the front and back of the load balancers:

Table 19.27. Application Ingress load balancer

Port	Back-end machines (pool members)	Internal	External	Description
------	----------------------------------	----------	----------	-------------

Port	Back-end machines (pool members)	Internal	External	Description
443	The machines that run the Ingress Controller pods, compute, or worker, by default.	X	X	HTTPS traffic
80	The machines that run the Ingress Controller pods, compute, or worker, by default.	X	X	HTTP traffic

**NOTE**

If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application Ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes.

19.3.4.8.1. Example load balancer configuration for user-provisioned clusters

This section provides an example API and application Ingress load balancer configuration that meets the load balancing requirements for user-provisioned clusters. The sample is an `/etc/haproxy/haproxy.cfg` configuration for an HAProxy load balancer. The example is not meant to provide advice for choosing one load balancing solution over another.

In the example, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

**NOTE**

If you are using HAProxy as a load balancer and SELinux is set to **enforcing**, you must ensure that the HAProxy service can bind to the configured TCP port by running **setsebool -P haproxy_connect_any=1**.

Example 19.6. Sample API and application Ingress load balancer configuration


```
global
  log      127.0.0.1 local2
  pidfile /var/run/haproxy.pid
  maxconn 4000
  daemon
defaults
  mode          http
  log           global
  option        dontlognull
  option http-server-close
  option        redispatch
  retries       3
  timeout http-request 10s
  timeout queue      1m
  timeout connect    10s
  timeout client     1m
```

```

timeout server      1m
timeout http-keep-alive 10s
timeout check      10s
maxconn            3000
listen api-server-6443 1
  bind *:6443
  mode tcp
  option httpchk GET /readyz HTTP/1.0
  option log-health-checks
  balance roundrobin
  server bootstrap bootstrap.ocp4.example.com:6443 verify none check check-ssl inter 10s fall 2
  rise 3 backup 2
  server master0 master0.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
  fall 2 rise 3
  server master1 master1.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
  fall 2 rise 3
  server master2 master2.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
  fall 2 rise 3
listen machine-config-server-22623 3
  bind *:22623
  mode tcp
  server bootstrap bootstrap.ocp4.example.com:22623 check inter 1s backup 4
  server master0 master0.ocp4.example.com:22623 check inter 1s
  server master1 master1.ocp4.example.com:22623 check inter 1s
  server master2 master2.ocp4.example.com:22623 check inter 1s
listen ingress-router-443 5
  bind *:443
  mode tcp
  balance source
  server compute0 compute0.ocp4.example.com:443 check inter 1s
  server compute1 compute1.ocp4.example.com:443 check inter 1s
listen ingress-router-80 6
  bind *:80
  mode tcp
  balance source
  server compute0 compute0.ocp4.example.com:80 check inter 1s
  server compute1 compute1.ocp4.example.com:80 check inter 1s

```

- 1** Port **6443** handles the Kubernetes API traffic and points to the control plane machines.
- 2** **4** The bootstrap entries must be in place before the OpenShift Container Platform cluster installation and they must be removed after the bootstrap process is complete.
- 3** Port **22623** handles the machine config server traffic and points to the control plane machines.
- 5** Port **443** handles the HTTPS traffic and points to the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.
- 6** Port **80** handles the HTTP traffic and points to the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.

**NOTE**

If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application Ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes.

TIP

If you are using HAProxy as a load balancer, you can check that the **haproxy** process is listening on ports **6443**, **22623**, **443**, and **80** by running **netstat -nltp** on the HAProxy node.

19.3.5. Preparing the user-provisioned infrastructure

Before you install OpenShift Container Platform on user-provisioned infrastructure, you must prepare the underlying infrastructure.

This section provides details about the high-level steps required to set up your cluster infrastructure in preparation for an OpenShift Container Platform installation. This includes configuring IP networking and network connectivity for your cluster nodes, enabling the required ports through your firewall, and setting up the required DNS and load balancing infrastructure.

After preparation, your cluster infrastructure must meet the requirements outlined in the *Requirements for a cluster with user-provisioned infrastructure* section.

Prerequisites

- You have reviewed the [OpenShift Container Platform 4.x Tested Integrations](#) page.
- You have reviewed the infrastructure requirements detailed in the *Requirements for a cluster with user-provisioned infrastructure* section.

Procedure

1. If you are using DHCP to provide the IP networking configuration to your cluster nodes, configure your DHCP service.
 - a. Add persistent IP addresses for the nodes to your DHCP server configuration. In your configuration, match the MAC address of the relevant network interface to the intended IP address for each node.
 - b. When you use DHCP to configure IP addressing for the cluster machines, the machines also obtain the DNS server information through DHCP. Define the persistent DNS server address that is used by the cluster nodes through your DHCP server configuration.

**NOTE**

If you are not using a DHCP service, you must provide the IP networking configuration and the address of the DNS server to the nodes at RHCOS install time. These can be passed as boot arguments if you are installing from an ISO image. See the *Installing RHCOS and starting the OpenShift Container Platform bootstrap process* section for more information about static IP provisioning and advanced networking options.

- c. Define the hostnames of your cluster nodes in your DHCP server configuration. See the *Setting the cluster node hostnames through DHCP* section for details about hostname considerations.

**NOTE**

If you are not using a DHCP service, the cluster nodes obtain their hostname through a reverse DNS lookup.

2. Ensure that your network infrastructure provides the required network connectivity between the cluster components. See the *Networking requirements for user-provisioned infrastructure* section for details about the requirements.
3. Configure your firewall to enable the ports required for the OpenShift Container Platform cluster components to communicate. See *Networking requirements for user-provisioned infrastructure* section for details about the ports that are required.

**IMPORTANT**

By default, port **1936** is accessible for an OpenShift Container Platform cluster, because each control plane node needs access to this port.

Avoid using the Ingress load balancer to expose this port, because doing so might result in the exposure of sensitive information, such as statistics and metrics, related to Ingress Controllers.


4. Setup the required DNS infrastructure for your cluster.
 - a. Configure DNS name resolution for the Kubernetes API, the application wildcard, the bootstrap machine, the control plane machines, and the compute machines.
 - b. Configure reverse DNS resolution for the Kubernetes API, the bootstrap machine, the control plane machines, and the compute machines.
See the *User-provisioned DNS requirements* section for more information about the OpenShift Container Platform DNS requirements.
5. Validate your DNS configuration.
 - a. From your installation node, run DNS lookups against the record names of the Kubernetes API, the wildcard routes, and the cluster nodes. Validate that the IP addresses in the responses correspond to the correct components.
 - b. From your installation node, run reverse DNS lookups against the IP addresses of the load balancer and the cluster nodes. Validate that the record names in the responses correspond to the correct components.
See the *Validating DNS resolution for user-provisioned infrastructure* section for detailed DNS validation steps.
6. Provision the required API and application ingress load balancing infrastructure. See the *Load balancing requirements for user-provisioned infrastructure* section for more information about the requirements.

**NOTE**

Some load balancing solutions require the DNS name resolution for the cluster nodes to be in place before the load balancing is initialized.

19.3.6. Validating DNS resolution for user-provisioned infrastructure

You can validate your DNS configuration before installing OpenShift Container Platform on user-provisioned infrastructure.

**IMPORTANT**

The validation steps detailed in this section must succeed before you install your cluster.

Prerequisites

- You have configured the required DNS records for your user-provisioned infrastructure.

Procedure

1. From your installation node, run DNS lookups against the record names of the Kubernetes API, the wildcard routes, and the cluster nodes. Validate that the IP addresses contained in the responses correspond to the correct components.
 - a. Perform a lookup against the Kubernetes API record name. Check that the result points to the IP address of the API load balancer:

```
$ dig +noall +answer @<nameserver_ip> api.<cluster_name>.<base_domain> 1
```

- 1** Replace **<nameserver_ip>** with the IP address of the nameserver, **<cluster_name>** with your cluster name, and **<base_domain>** with your base domain name.

Example output

```
api.ocp4.example.com. 604800 IN A 192.168.1.5
```

- b. Perform a lookup against the Kubernetes internal API record name. Check that the result points to the IP address of the API load balancer:

```
$ dig +noall +answer @<nameserver_ip> api-int.<cluster_name>.<base_domain>
```

Example output

```
api-int.ocp4.example.com. 604800 IN A 192.168.1.5
```

- c. Test an example ***.apps.<cluster_name>.<base_domain>** DNS wildcard lookup. All of the application wildcard lookups must resolve to the IP address of the application ingress load balancer:

```
$ dig +noall +answer @<nameserver_ip> random.apps.<cluster_name>.<base_domain>
```

Example output

```
random.apps.ocp4.example.com. 604800 IN A 192.168.1.5
```



NOTE

In the example outputs, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

You can replace **random** with another wildcard value. For example, you can query the route to the OpenShift Container Platform console:

```
$ dig +noall +answer @<nameserver_ip> console-openshift-console.apps.<cluster_name>.<base_domain>
```

Example output

```
console-openshift-console.apps.ocp4.example.com. 604800 IN A 192.168.1.5
```

- d. Run a lookup against the bootstrap DNS record name. Check that the result points to the IP address of the bootstrap node:

```
$ dig +noall +answer @<nameserver_ip> bootstrap.<cluster_name>.<base_domain>
```

Example output

```
bootstrap.ocp4.example.com. 604800 IN A 192.168.1.96
```

- e. Use this method to perform lookups against the DNS record names for the control plane and compute nodes. Check that the results correspond to the IP addresses of each node.
2. From your installation node, run reverse DNS lookups against the IP addresses of the load balancer and the cluster nodes. Validate that the record names contained in the responses correspond to the correct components.
 - a. Perform a reverse lookup against the IP address of the API load balancer. Check that the response includes the record names for the Kubernetes API and the Kubernetes internal API:

```
$ dig +noall +answer @<nameserver_ip> -x 192.168.1.5
```

Example output

```
5.1.168.192.in-addr.arpa. 604800 IN PTR api-int.ocp4.example.com. 1  
5.1.168.192.in-addr.arpa. 604800 IN PTR api.ocp4.example.com. 2
```

1 Provides the record name for the Kubernetes internal API.

2 Provides the record name for the Kubernetes API.

**NOTE**

A PTR record is not required for the OpenShift Container Platform application wildcard. No validation step is needed for reverse DNS resolution against the IP address of the application ingress load balancer.

- b. Perform a reverse lookup against the IP address of the bootstrap node. Check that the result points to the DNS record name of the bootstrap node:

```
$ dig +noall +answer @<nameserver_ip> -x 192.168.1.96
```

Example output

```
96.1.168.192.in-addr.arpa. 604800 IN PTR bootstrap.ocp4.example.com.
```


- c. Use this method to perform reverse lookups against the IP addresses for the control plane and compute nodes. Check that the results correspond to the DNS record names of each node.

19.3.7. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.

**IMPORTANT**

Do not skip this procedure in production environments, where disaster recovery and debugging is required.

**NOTE**


You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.



NOTE

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.


2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.



NOTE

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```



NOTE

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1 Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
-
```



```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

19.3.8. Manually creating the installation configuration file

Installing the cluster requires that you manually create the installation configuration file.


Prerequisites

- You have an SSH public key on your local machine to provide to the installation program. The key will be used for SSH authentication onto your cluster nodes for debugging and disaster recovery.
- You have obtained the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create an installation directory to store your required installation assets in:

```
$ mkdir <installation_directory>
```



IMPORTANT

You must create a directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.


2. Customize the sample **install-config.yaml** file template that is provided and save it in the **<installation_directory>**.



NOTE

You must name this configuration file **install-config.yaml**.

3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the next step of the installation process. You must back it up now.

Additional resources

- [Installation configuration parameters for IBM Power®](#)

- 4 You must set this value to **0** when you install OpenShift Container Platform on user-provisioned infrastructure. In installer-provisioned installations, the parameter controls the number of compute



NOTE

If you are installing a three-node cluster, do not deploy any compute machines when you install the Red Hat Enterprise Linux CoreOS (RHCOS) machines.

- 7 The number of control plane machines that you add to the cluster. Because the cluster uses these values as the number of etcd endpoints in the cluster, the value must match the number of control plane machines that you deploy.
- 8 The cluster name that you specified in your DNS records.
- 9 A block of IP addresses from which pod IP addresses are allocated. This block must not overlap with existing physical networks. These IP addresses are used for the pod network. If you need to access the pods from an external network, you must configure load balancers and routers to manage the traffic.



NOTE

Class E CIDR range is reserved for a future use. To use the Class E CIDR range, you must ensure your networking environment accepts the IP addresses within the Class E CIDR range.

- 10 The subnet prefix length to assign to each individual node. For example, if **hostPrefix** is set to **23**, then each node is assigned a **/23** subnet out of the given **cidr**, which allows for 510 ($2^{(32 - 23)} - 2$) pod IP addresses. If you are required to provide access to nodes from an external network, configure load balancers and routers to manage the traffic.
- 11 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 12 The IP address pool to use for service IP addresses. You can enter only one IP address pool. This block must not overlap with existing physical networks. If you need to access the services from an external network, configure load balancers and routers to manage the traffic.
- 13 You must set the platform to **none**. You cannot provide additional platform configuration variables for IBM Power® infrastructure.



IMPORTANT

Clusters that are installed with the platform type **none** are unable to use some features, such as managing compute machines with the Machine API. This limitation applies even if the compute machines that are attached to the cluster are installed on a platform that would normally support the feature. This parameter cannot be changed after installation.

- 14 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.



IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 15 For **<local_registry>**, specify the registry domain name, and optionally the port, that your mirror registry uses to serve content. For example, **registry.example.com** or **registry.example.com:5000**. For **<credentials>**, specify the base64-encoded user name and password for your mirror registry.
- 16 The SSH public key for the **core** user in Red Hat Enterprise Linux CoreOS (RHCOS).



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- 17 Provide the contents of the certificate file that you used for your mirror registry.
- 18 Provide the **imageContentSources** section according to the output of the command that you used to mirror the repository.



IMPORTANT

- When using the **oc adm release mirror** command, use the output from the **imageContentSources** section.
- When using **oc mirror** command, use the **repositoryDigestMirrors** section of the **ImageContentSourcePolicy** file that results from running the command.
- **ImageContentSourcePolicy** is deprecated. For more information see *Configuring image registry repository mirroring*.

19.3.8.2. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.

- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.


For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

- Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- A proxy URL to use for creating HTTPS connections outside the cluster.
- A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations.
- If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.

**NOTE**

The installation program does not support the proxy **readinessEndpoints** field.

**NOTE**

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.

**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

19.3.8.3. Configuring a three-node cluster

Optionally, you can deploy zero compute machines in a bare metal cluster that consists of three control plane machines only. This provides smaller, more resource efficient clusters for cluster administrators and developers to use for testing, development, and production.

In three-node OpenShift Container Platform environments, the three control plane machines are schedulable, which means that your application workloads are scheduled to run on them.


Prerequisites

- You have an existing **install-config.yaml** file.

Procedure

- Ensure that the number of compute replicas is set to **0** in your **install-config.yaml** file, as shown in the following **compute** stanza:

```
compute:
- name: worker
  platform: {}
  replicas: 0
```

**NOTE**

You must set the value of the **replicas** parameter for the compute machines to **0** when you install OpenShift Container Platform on user-provisioned infrastructure, regardless of the number of compute machines you are deploying. In installer-provisioned installations, the parameter controls the number of compute machines that the cluster creates and manages for you. This does not apply to user-provisioned installations, where the compute machines are deployed manually.

For three-node cluster installations, follow these next steps:

- If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes. See the *Load balancing requirements for user-provisioned infrastructure* section for more information.
- When you create the Kubernetes manifest files in the following procedure, ensure that the **mastersSchedulable** parameter in the `<installation_directory>/manifests/cluster-scheduler-02-config.yml` file is set to **true**. This enables your application workloads to run on the control plane nodes.
- Do not deploy any compute nodes when you create the Red Hat Enterprise Linux CoreOS (RHCOS) machines.

19.3.9. Cluster Network Operator configuration

The configuration for the cluster network is specified as part of the Cluster Network Operator (CNO) configuration and stored in a custom resource (CR) object that is named **cluster**. The CR specifies the fields for the **Network** API in the **operator.openshift.io** API group.

The CNO configuration inherits the following fields during cluster installation from the **Network** API in the **Network.config.openshift.io** API group:

clusterNetwork

IP address pools from which pod IP addresses are allocated.

serviceNetwork

IP address pool for services.

defaultNetwork.type

Cluster network plugin. **OVNKubernetes** is the only supported plugin during installation.

You can specify the cluster network plugin configuration for your cluster by setting the fields for the **defaultNetwork** object in the CNO object named **cluster**.

19.3.9.1. Cluster Network Operator configuration object

The fields for the Cluster Network Operator (CNO) are described in the following table:

Table 19.28. Cluster Network Operator configuration object


Field	Type	Description
metadata.name	string	The name of the CNO object. This name is always cluster .
spec.clusterNetwork	array	A list specifying the blocks of IP addresses from which pod IP addresses are allocated and the subnet prefix length assigned to each individual node in the cluster. For example: <pre>spec: clusterNetwork: - cidr: 10.128.0.0/19 hostPrefix: 23 - cidr: 10.128.32.0/19 hostPrefix: 23</pre>
spec.serviceNetwork	array	A block of IP addresses for services. The OpenShift SDN and OVN-Kubernetes network plugins support only a single IP address block for the service network. For example: <pre>spec: serviceNetwork: - 172.30.0.0/14</pre> <p>You can customize this field only in the install-config.yaml file before you create the manifests. The value is read-only in the manifest file.</p>
spec.defaultNetwork	object	Configures the network plugin for the cluster network.
spec.kubeProxy Config	object	The fields for this object specify the kube-proxy configuration. If you are using the OVN-Kubernetes cluster network plugin, the kube-proxy configuration has no effect.

defaultNetwork object configuration

The values for the **defaultNetwork** object are defined in the following table:

Table 19.29. **defaultNetwork** object

Field	Type	Description
-------	------	-------------

Field	Type	Description
type	string	<p>OVNKubernetes. The Red Hat OpenShift Networking network plugin is selected during installation. This value cannot be changed after cluster installation.</p> <div style="display: flex; align-items: flex-start;">  <p>NOTE</p> <p>OpenShift Container Platform uses the OVN-Kubernetes network plugin by default. OpenShift SDN is no longer available as an installation choice for new clusters.</p> </div>
ovnKubernetesConfig	object	This object is only valid for the OVN-Kubernetes network plugin.

Configuration for the OVN-Kubernetes network plugin

The following table describes the configuration fields for the OVN-Kubernetes network plugin:

Table 19.30. ovnKubernetesConfig object

Field	Type	Description
mtu	integer	<p>The maximum transmission unit (MTU) for the Geneve (Generic Network Virtualization Encapsulation) overlay network. This is detected automatically based on the MTU of the primary network interface. You do not normally need to override the detected MTU.</p> <p>If the auto-detected value is not what you expect it to be, confirm that the MTU on the primary network interface on your nodes is correct. You cannot use this option to change the MTU value of the primary network interface on the nodes.</p> <p>If your cluster requires different MTU values for different nodes, you must set this value to 100 less than the lowest MTU value in your cluster. For example, if some nodes in your cluster have an MTU of 9001, and some have an MTU of 1500, you must set this value to 1400.</p>
genevePort	integer	The port to use for all Geneve packets. The default value is 6081 . This value cannot be changed after cluster installation.
ipsecConfig	object	Specify a configuration object for customizing the IPsec configuration.
ipv4	object	Specifies a configuration object for IPv4 settings.
ipv6	object	Specifies a configuration object for IPv6 settings.


Field	Type	Description
policyAuditConfig	object	Specify a configuration object for customizing network policy audit logging. If unset, the defaults audit log settings are used.
gatewayConfig	object	Optional: Specify a configuration object for customizing how egress traffic is sent to the node gateway.  <p>NOTE</p> <p>While migrating egress traffic, you can expect some disruption to workloads and service traffic until the Cluster Network Operator (CNO) successfully rolls out the changes.</p>

Table 19.31. `ovnKubernetesConfig.ipv4` object

Field	Type	Description
internalTransitSwitchSubnet	string	If your existing network infrastructure overlaps with the 100.88.0.0/16 IPv4 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. The subnet for the distributed transit switch that enables east-west traffic. This subnet cannot overlap with any other subnets used by OVN-Kubernetes or on the host itself. It must be large enough to accommodate one IP address per node in your cluster. The default value is 100.88.0.0/16 .
internalJoinSubnet	string	If your existing network infrastructure overlaps with the 100.64.0.0/16 IPv4 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. You must ensure that the IP address range does not overlap with any other subnet used by your OpenShift Container Platform installation. The IP address range must be larger than the maximum number of nodes that can be added to the cluster. For example, if the clusterNetwork.cidr value is 10.128.0.0/14 and the clusterNetwork.hostPrefix value is /23 , then the maximum number of nodes is 2⁽²³⁻¹⁴⁾=512 . The default value is 100.64.0.0/16 .

Table 19.32. `ovnKubernetesConfig.ipv6` object

Field	Type	Description
-------	------	-------------

Field	Type	Description
internalTransitSwitchSubnet	string	<p>If your existing network infrastructure overlaps with the fd97::/64 IPv6 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. The subnet for the distributed transit switch that enables east-west traffic. This subnet cannot overlap with any other subnets used by OVN-Kubernetes or on the host itself. It must be large enough to accommodate one IP address per node in your cluster.</p> <p>The default value is fd97::/64.</p>
internalJoinSubnet	string	<p>If your existing network infrastructure overlaps with the fd98::/64 IPv6 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. You must ensure that the IP address range does not overlap with any other subnet used by your OpenShift Container Platform installation. The IP address range must be larger than the maximum number of nodes that can be added to the cluster.</p> <p>The default value is fd98::/64.</p>

Table 19.33. **policyAuditConfig** object

Field	Type	Description
rateLimit	integer	The maximum number of messages to generate every second per node. The default value is 20 messages per second.
maxFileSize	integer	The maximum size for the audit log in bytes. The default value is 50000000 or 50 MB.
maxLogFiles	integer	The maximum number of log files that are retained.
destination	string	<p>One of the following additional audit log targets:</p> <p>libc The libc syslog() function of the journald process on the host.</p> <p>udp:<host>:<port> A syslog server. Replace <host>:<port> with the host and port of the syslog server.</p> <p>unix:<file> A Unix Domain Socket file specified by <file>.</p> <p>null Do not send the audit logs to any additional target.</p>
syslogFacility	string	The syslog facility, such as kern , as defined by RFC5424. The default value is local0 .

Table 19.34. gatewayConfig object

Field	Type	Description
routingViaHost	boolean	<p>Set this field to true to send egress traffic from pods to the host networking stack. For highly-specialized installations and applications that rely on manually configured routes in the kernel routing table, you might want to route egress traffic to the host networking stack. By default, egress traffic is processed in OVN to exit the cluster and is not affected by specialized routes in the kernel routing table. The default value is false.</p> <p>This field has an interaction with the Open vSwitch hardware offloading feature. If you set this field to true, you do not receive the performance benefits of the offloading because egress traffic is processed by the host networking stack.</p>
ipForwarding	object	You can control IP forwarding for all traffic on OVN-Kubernetes managed interfaces by using the ipForwarding specification in the Network resource. Specify Restricted to only allow IP forwarding for Kubernetes related traffic. Specify Global to allow forwarding of all IP traffic. For new installations, the default is Restricted . For updates to OpenShift Container Platform 4.14 or later, the default is Global .
ipv4	object	Optional: Specify an object to configure the internal OVN-Kubernetes masquerade address for host to service traffic for IPv4 addresses.
ipv6	object	Optional: Specify an object to configure the internal OVN-Kubernetes masquerade address for host to service traffic for IPv6 addresses.

Table 19.35. gatewayConfig.ipv4 object

Field	Type	Description
internalMasqueradeSubnet	string	The masquerade IPv4 addresses that are used internally to enable host to service traffic. The host is configured with these IP addresses as well as the shared gateway bridge interface. The default value is 169.254.169.0/29 .

Table 19.36. gatewayConfig.ipv6 object

Field	Type	Description
internalMasqueradeSubnet	string	The masquerade IPv6 addresses that are used internally to enable host to service traffic. The host is configured with these IP addresses as well as the shared gateway bridge interface. The default value is fd69::/125 .

Table 19.37. `ipsecConfig` object


Field	Type	Description
<code>mode</code>	<code>string</code>	<p>Specifies the behavior of the IPsec implementation. Must be one of the following values:</p> <ul style="list-style-type: none"> ● Disabled: IPsec is not enabled on cluster nodes. ● External: IPsec is enabled for network traffic with external hosts. ● Full: IPsec is enabled for pod traffic and network traffic with external hosts.

Example OVN-Kubernetes configuration with IPsec enabled

```

defaultNetwork:
  type: OVNKubernetes
  ovnKubernetesConfig:
    mtu: 1400
    genevePort: 6081
    ipsecConfig:
      mode: Full

```




IMPORTANT

Using OVNKubernetes can lead to a stack exhaustion problem on IBM Power®.

`kubeProxyConfig` object configuration (OpenShiftSDN container network interface only)

The values for the `kubeProxyConfig` object are defined in the following table:

Table 19.38. `kubeProxyConfig` object

Field	Type	Description
<code>iptablesSyncPeriod</code>	<code>string</code>	<p>The refresh period for <code>iptables</code> rules. The default value is 30s. Valid suffixes include s, m, and h and are described in the Go time package documentation.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p>NOTE</p> <p>Because of performance improvements introduced in OpenShift Container Platform 4.3 and greater, adjusting the <code>iptablesSyncPeriod</code> parameter is no longer necessary.</p> </div> </div>

Field	Type	Description
proxyArguments.iptables-min-sync-period	array	<p>The minimum duration before refreshing iptables rules. This field ensures that the refresh does not happen too frequently. Valid suffixes include s, m, and h and are described in the Go time package. The default value is:</p> <pre>kubeProxyConfig: proxyArguments: iptables-min-sync-period: - 0s</pre>

19.3.10. Creating the Kubernetes manifest and Ignition config files

Because you must modify some cluster definition files and manually start the cluster machines, you must generate the Kubernetes manifest and Ignition config files that the cluster needs to configure the machines.

The installation configuration file transforms into the Kubernetes manifests. The manifests wrap into the Ignition configuration files, which are later used to configure the cluster machines.



IMPORTANT

- The Ignition config files that the OpenShift Container Platform installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.



NOTE

The installation program that generates the manifest and Ignition files is architecture specific and can be obtained from the [client image mirror](#). The Linux version of the installation program (without an architecture postfix) runs on ppc64le only. This installer program is also available as a Mac OS version.

Prerequisites

- You obtained the OpenShift Container Platform installation program. For a restricted network installation, these files are on your mirror host.
- You created the **install-config.yaml** installation configuration file.

Procedure

1. Change to the directory that contains the OpenShift Container Platform installation program and generate the Kubernetes manifests for the cluster:


```
$ ./openshift-install create manifests --dir <installation_directory> 1
```

- 1 For **<installation_directory>**, specify the installation directory that contains the **install-config.yaml** file you created.



WARNING

If you are installing a three-node cluster, skip the following step to allow the control plane nodes to be schedulable.



IMPORTANT

When you configure control plane nodes from the default unschedulable to schedulable, additional subscriptions are required. This is because control plane nodes then become compute nodes.

2. Check that the **mastersSchedulable** parameter in the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** Kubernetes manifest file is set to **false**. This setting prevents pods from being scheduled on the control plane machines:
 - a. Open the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** file.
 - b. Locate the **mastersSchedulable** parameter and ensure that it is set to **false**.
 - c. Save and exit the file.
3. To create the Ignition configuration files, run the following command from the directory that contains the installation program:

```
$ ./openshift-install create ignition-configs --dir <installation_directory> 1
```

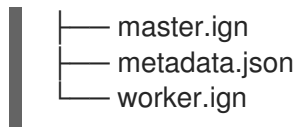
- 1 For **<installation_directory>**, specify the same installation directory.

Ignition config files are created for the bootstrap, control plane, and compute nodes in the installation directory. The **kubeadmin-password** and **kubeconfig** files are created in the **./<installation_directory>/auth** directory:

```

.
├── auth
│   ├── kubeadmin-password
│   └── kubeconfig
└── bootstrap.ign

```



19.3.11. Installing RHCOS and starting the OpenShift Container Platform bootstrap process

To install OpenShift Container Platform on IBM Power® infrastructure that you provision, you must install Red Hat Enterprise Linux CoreOS (RHCOS) on the machines. When you install RHCOS, you must provide the Ignition config file that was generated by the OpenShift Container Platform installation program for the type of machine you are installing. If you have configured suitable networking, DNS, and load balancing infrastructure, the OpenShift Container Platform bootstrap process begins automatically after the RHCOS machines have rebooted.

Follow either the steps to use an ISO image or network PXE booting to install RHCOS on the machines.

19.3.11.1. Installing RHCOS by using an ISO image

You can use an ISO image to install RHCOS on the machines.

Prerequisites

- You have created the Ignition config files for your cluster.
- You have configured suitable network, DNS and load balancing infrastructure.
- You have an HTTP server that can be accessed from your computer, and from the machines that you create.
- You have reviewed the *Advanced RHCOS installation configuration* section for different ways to configure features, such as networking and disk partitioning.


Procedure

1. Obtain the SHA512 digest for each of your Ignition config files. For example, you can use the following on a system running Linux to get the SHA512 digest for your **bootstrap.ign** Ignition config file:

```
$ sha512sum <installation_directory>/bootstrap.ign
```

The digests are provided to the **coreos-installer** in a later step to validate the authenticity of the Ignition config files on the cluster nodes.

2. Upload the bootstrap, control plane, and compute node Ignition config files that the installation program created to your HTTP server. Note the URLs of these files.



IMPORTANT

You can add or change configuration settings in your Ignition configs before saving them to your HTTP server. If you plan to add more compute machines to your cluster after you finish installation, do not delete these files.

3. From the installation host, validate that the Ignition config files are available on the URLs. The following example gets the Ignition config file for the bootstrap node:


```
$ curl -k http://<HTTP_server>/bootstrap.ign 1
```

Example output

```
% Total % Received % Xferd Average Speed Time Time Time Current
          Dload Upload Total Spent Left Speed
0 0 0 0 0 0 0 0 --:--:-- --:--:-- --:--:-- 0{"ignition":
{"version":"3.2.0"},"passwd":{"users":[{"name":"core","sshAuthorizedKeys":["ssh-rsa...
```


Replace **bootstrap.ign** with **master.ign** or **worker.ign** in the command to validate that the Ignition config files for the control plane and compute nodes are also available.

- Although it is possible to obtain the RHCOS images that are required for your preferred method of installing operating system instances from the [RHCOS image mirror](#) page, the recommended way to obtain the correct version of your RHCOS images are from the output of **openshift-install** command:

```
$ openshift-install coreos print-stream-json | grep '\.iso[^\.]'
```

Example output

```
"location": "<url>/art/storage/releases/rhcos-4.16-aarch64/<release>/aarch64/rhcos-
<release>-live.aarch64.iso",
"location": "<url>/art/storage/releases/rhcos-4.16-ppc64le/<release>/ppc64le/rhcos-
<release>-live.ppc64le.iso",
"location": "<url>/art/storage/releases/rhcos-4.16-s390x/<release>/s390x/rhcos-<release>-
live.s390x.iso",
"location": "<url>/art/storage/releases/rhcos-4.16/<release>/x86_64/rhcos-<release>-
live.x86_64.iso",
```



IMPORTANT

The RHCOS images might not change with every release of OpenShift Container Platform. You must download images with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Use the image versions that match your OpenShift Container Platform version if they are available. Use only ISO images for this procedure. RHCOS qcow2 images are not supported for this installation type.

ISO file names resemble the following example:

rhcos-<version>-live.<architecture>.iso

- Use the ISO to start the RHCOS installation. Use one of the following installation options:
 - Burn the ISO image to a disk and boot it directly.
 - Use ISO redirection by using a lights-out management (LOM) interface.
- Boot the RHCOS ISO image without specifying any options or interrupting the live boot sequence. Wait for the installer to boot into a shell prompt in the RHCOS live environment.

**NOTE**

It is possible to interrupt the RHCOS installation boot process to add kernel arguments. However, for this ISO procedure you should use the **coreos-installer** command as outlined in the following steps, instead of adding kernel arguments.

7. Run the **coreos-installer** command and specify the options that meet your installation requirements. At a minimum, you must specify the URL that points to the Ignition config file for the node type, and the device that you are installing to:

```
$ sudo coreos-installer install --ignition-url=http://<HTTP_server>/<node_type>.ign <device>
--ignition-hash=sha512-<digest> 1 2
```

- 1** You must run the **coreos-installer** command by using **sudo**, because the **core** user does not have the required root privileges to perform the installation.
- 2** The **--ignition-hash** option is required when the Ignition config file is obtained through an HTTP URL to validate the authenticity of the Ignition config file on the cluster node. **<digest>** is the Ignition config file SHA512 digest obtained in a preceding step.

**NOTE**

If you want to provide your Ignition config files through an HTTPS server that uses TLS, you can add the internal certificate authority (CA) to the system trust store before running **coreos-installer**.

The following example initializes a bootstrap node installation to the **/dev/sda** device. The Ignition config file for the bootstrap node is obtained from an HTTP web server with the IP address 192.168.1.2:

```
$ sudo coreos-installer install --ignition-
url=http://192.168.1.2:80/installation_directory/bootstrap.ign /dev/sda --ignition-hash=sha512-
a5a2d43879223273c9b60af66b44202a1d1248fc01cf156c46d4a79f552b6bad47bc8cc78ddf011
6e80c59d2ea9e32ba53bc807afbca581aa059311def2c3e3b
```

8. Monitor the progress of the RHCOS installation on the console of the machine.

**IMPORTANT**

Be sure that the installation is successful on each node before commencing with the OpenShift Container Platform installation. Observing the installation process can also help to determine the cause of RHCOS installation issues that might arise.

9. After RHCOS installs, you must reboot the system. During the system reboot, it applies the Ignition config file that you specified.
10. Check the console output to verify that Ignition ran.

Example command

Ignition: ran on 2022/03/14 14:48:33 UTC (this boot)
 Ignition: user-provided config was applied

11. Continue to create the other machines for your cluster.



IMPORTANT

You must create the bootstrap and control plane machines at this time. If the control plane machines are not made schedulable, also create at least two compute machines before you install OpenShift Container Platform.

If the required network, DNS, and load balancer infrastructure are in place, the OpenShift Container Platform bootstrap process begins automatically after the RHCOS nodes have rebooted.



NOTE


RHCOS nodes do not include a default password for the **core** user. You can access the nodes by running **ssh core@<node>.<cluster_name>.<base_domain>** as a user with access to the SSH private key that is paired to the public key that you specified in your **install_config.yaml** file. OpenShift Container Platform 4 cluster nodes running RHCOS are immutable and rely on Operators to apply cluster changes. Accessing cluster nodes by using SSH is not recommended. However, when investigating installation issues, if the OpenShift Container Platform API is not available, or the kubelet is not properly functioning on a target node, SSH access might be required for debugging or disaster recovery.

19.3.11.1.1. Advanced RHCOS installation reference

This section illustrates the networking configuration and other advanced options that allow you to modify the Red Hat Enterprise Linux CoreOS (RHCOS) manual installation process. The following tables describe the kernel arguments and command-line options you can use with the RHCOS live installer and the **coreos-installer** command.

19.3.11.1.1.1. Networking and bonding options for ISO installations

If you install RHCOS from an ISO image, you can add kernel arguments manually when you boot the image to configure networking for a node. If no networking arguments are specified, DHCP is activated in the initramfs when RHCOS detects that networking is required to fetch the Ignition config file.



IMPORTANT

When adding networking arguments manually, you must also add the **rd.neednet=1** kernel argument to bring the network up in the initramfs.

The following information provides examples for configuring networking and bonding on your RHCOS nodes for ISO installations. The examples describe how to use the **ip=**, **nameserver=**, and **bond=** kernel arguments.

**NOTE**

Ordering is important when adding the kernel arguments: **ip=**, **nameserver=**, and then **bond=**.

The networking options are passed to the **dracut** tool during system boot. For more information about the networking options supported by **dracut**, see the [dracut.cmdline manual page](#).

The following examples are the networking options for ISO installation.

Configuring DHCP or static IP addresses

To configure an IP address, either use DHCP (**ip=dhcp**) or set an individual static IP address (**ip=<host_ip>**). If setting a static IP, you must then identify the DNS server IP address (**nameserver=<dns_ip>**) on each node. The following example sets:

- The node's IP address to **10.10.10.2**
- The gateway address to **10.10.10.254**
- The netmask to **255.255.255.0**
- The hostname to **core0.example.com**
- The DNS server address to **4.4.4.41**
- The auto-configuration value to **none**. No auto-configuration is required when IP networking is configured statically.

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none
nameserver=4.4.4.41
```

**NOTE**

When you use DHCP to configure IP addressing for the RHCOS machines, the machines also obtain the DNS server information through DHCP. For DHCP-based deployments, you can define the DNS server address that is used by the RHCOS nodes through your DHCP server configuration.

Configuring an IP address without a static hostname

You can configure an IP address without assigning a static hostname. If a static hostname is not set by the user, it will be picked up and automatically set by a reverse DNS lookup. To configure an IP address without a static hostname refer to the following example:

- The node's IP address to **10.10.10.2**
- The gateway address to **10.10.10.254**
- The netmask to **255.255.255.0**
- The DNS server address to **4.4.4.41**
- The auto-configuration value to **none**. No auto-configuration is required when IP networking is configured statically.

```
ip=10.10.10.2::10.10.10.254:255.255.255.0::enp1s0:none
nameserver=4.4.4.41
```

Specifying multiple network interfaces

You can specify multiple network interfaces by setting multiple **ip=** entries.

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none
ip=10.10.10.3::10.10.10.254:255.255.255.0:core0.example.com:enp2s0:none
```

Configuring default gateway and route

Optional: You can configure routes to additional networks by setting an **rd.route=** value.



NOTE

When you configure one or multiple networks, one default gateway is required. If the additional network gateway is different from the primary network gateway, the default gateway must be the primary network gateway.

- Run the following command to configure the default gateway:

```
ip=::10.10.10.254:::
```

- Enter the following command to configure the route for the additional network:

```
rd.route=20.20.20.0/24:20.20.20.254:enp2s0
```

Disabling DHCP on a single interface

You can disable DHCP on a single interface, such as when there are two or more network interfaces and only one interface is being used. In the example, the **enp1s0** interface has a static networking configuration and DHCP is disabled for **enp2s0**, which is not used:

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none
ip=::::core0.example.com:enp2s0:none
```

Combining DHCP and static IP configurations

You can combine DHCP and static IP configurations on systems with multiple network interfaces, for example:

```
ip=enp1s0:dhcp
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp2s0:none
```

Configuring VLANs on individual interfaces

Optional: You can configure VLANs on individual interfaces by using the **vlan=** parameter.

- To configure a VLAN on a network interface and use a static IP address, run the following command:

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp2s0.100:none
vlan=enp2s0.100:enp2s0
```

- To configure a VLAN on a network interface and to use DHCP, run the following command:

```
ip=enp2s0.100:dhcp
vlan=enp2s0.100:enp2s0
```

Providing multiple DNS servers

You can provide multiple DNS servers by adding a **nameserver=** entry for each server, for example:

```
nameserver=1.1.1.1
nameserver=8.8.8.8
```

Bonding multiple network interfaces to a single interface

Optional: You can bond multiple network interfaces to a single interface by using the **bond=** option.

Refer to the following examples:

- The syntax for configuring a bonded interface is: **bond=<name>[:<network_interfaces>][:options]**
<name> is the bonding device name (**bond0**), **<network_interfaces>** represents a comma-separated list of physical (ethernet) interfaces (**em1,em2**), and *options* is a comma-separated list of bonding options. Enter **modinfo bonding** to see available options.
- When you create a bonded interface using **bond=**, you must specify how the IP address is assigned and other information for the bonded interface.
 - To configure the bonded interface to use DHCP, set the bond's IP address to **dhcp**. For example:

```
bond=bond0:em1,em2:mode=active-backup
ip=bond0:dhcp
```

- To configure the bonded interface to use a static IP address, enter the specific IP address you want and related information. For example:

```
bond=bond0:em1,em2:mode=active-backup
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:bond0:none
```

Bonding multiple SR-IOV network interfaces to a dual port NIC interface



IMPORTANT

Support for Day 1 operations associated with enabling NIC partitioning for SR-IOV devices is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

Optional: You can bond multiple SR-IOV network interfaces to a dual port NIC interface by using the **bond=** option.

On each node, you must perform the following tasks:

1. Create the SR-IOV virtual functions (VFs) following the guidance in [Managing SR-IOV devices](#). Follow the procedure in the "Attaching SR-IOV networking devices to virtual machines" section.
2. Create the bond, attach the desired VFs to the bond and set the bond link state up following the guidance in [Configuring network bonding](#). Follow any of the described procedures to create the bond.

The following examples illustrate the syntax you must use:

- The syntax for configuring a bonded interface is **bond=<name>[:<network_interfaces>][:options]**. **<name>** is the bonding device name (**bond0**), **<network_interfaces>** represents the virtual functions (VFs) by their known name in the kernel and shown in the output of the **ip link** command (**eno1f0**, **eno2f0**), and *options* is a comma-separated list of bonding options. Enter **modinfo bonding** to see available options.
- When you create a bonded interface using **bond=**, you must specify how the IP address is assigned and other information for the bonded interface.
 - To configure the bonded interface to use DHCP, set the bond's IP address to **dhcp**. For example:

```
bond=bond0:eno1f0,eno2f0:mode=active-backup
ip=bond0:dhcp
```

- To configure the bonded interface to use a static IP address, enter the specific IP address you want and related information. For example:

```
bond=bond0:eno1f0,eno2f0:mode=active-backup
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:bond0:none
```

19.3.11.2. Installing RHCOS by using PXE booting


You can use PXE booting to install RHCOS on the machines.

Prerequisites

- You have created the Ignition config files for your cluster.
- You have configured suitable network, DNS and load balancing infrastructure.
- You have configured suitable PXE infrastructure.
- You have an HTTP server that can be accessed from your computer, and from the machines that you create.
- You have reviewed the *Advanced RHCOS installation configuration* section for different ways to configure features, such as networking and disk partitioning.

Procedure

1. Upload the bootstrap, control plane, and compute node Ignition config files that the installation program created to your HTTP server. Note the URLs of these files.



IMPORTANT

You can add or change configuration settings in your Ignition configs before saving them to your HTTP server. If you plan to add more compute machines to your cluster after you finish installation, do not delete these files.

- From the installation host, validate that the Ignition config files are available on the URLs. The following example gets the Ignition config file for the bootstrap node:

```
$ curl -k http://<HTTP_server>/bootstrap.ign 1
```

Example output

```
% Total    % Received % Xferd Average Speed   Time    Time     Time  Current
           Dload  Upload   Total   Spent    Left  Speed
  0   0   0   0   0   0   0   0   0  --:--:-- --:--:-- --:--:--    0{"ignition":
{"version":"3.2.0"},"passwd":{"users":[{"name":"core","sshAuthorizedKeys":["ssh-rsa...
```

Replace **bootstrap.ign** with **master.ign** or **worker.ign** in the command to validate that the Ignition config files for the control plane and compute nodes are also available.

- Although it is possible to obtain the RHCOS **kernel**, **initramfs** and **rootfs** files that are required for your preferred method of installing operating system instances from the [RHCOS image mirror](#) page, the recommended way to obtain the correct version of your RHCOS files are from the output of **openshift-install** command:

```
$ openshift-install coreos print-stream-json | grep -Eo "https.*(kernel-|initramfs-|rootfs-)\w+(\.img)?"
```

Example output

```
"<url>/art/storage/releases/rhcos-4.16-aarch64/<release>/aarch64/rhcos-<release>-live-
kernel-aarch64"
"<url>/art/storage/releases/rhcos-4.16-aarch64/<release>/aarch64/rhcos-<release>-live-
initramfs.aarch64.img"
"<url>/art/storage/releases/rhcos-4.16-aarch64/<release>/aarch64/rhcos-<release>-live-
rootfs.aarch64.img"
"<url>/art/storage/releases/rhcos-4.16-ppc64le/49.84.202110081256-0/ppc64le/rhcos-
<release>-live-kernel-ppc64le"
"<url>/art/storage/releases/rhcos-4.16-ppc64le/<release>/ppc64le/rhcos-<release>-live-
initramfs.ppc64le.img"
"<url>/art/storage/releases/rhcos-4.16-ppc64le/<release>/ppc64le/rhcos-<release>-live-
rootfs.ppc64le.img"
"<url>/art/storage/releases/rhcos-4.16-s390x/<release>/s390x/rhcos-<release>-live-kernel-
s390x"
"<url>/art/storage/releases/rhcos-4.16-s390x/<release>/s390x/rhcos-<release>-live-
initramfs.s390x.img"
"<url>/art/storage/releases/rhcos-4.16-s390x/<release>/s390x/rhcos-<release>-live-
rootfs.s390x.img"
"<url>/art/storage/releases/rhcos-4.16/<release>/x86_64/rhcos-<release>-live-kernel-
x86_64"
"<url>/art/storage/releases/rhcos-4.16/<release>/x86_64/rhcos-<release>-live-
```



```
initramfs.x86_64.img"
"<url>/art/storage/releases/rhcos-4.16/<release>/x86_64/rhcos-<release>-live-
rootfs.x86_64.img"
```



IMPORTANT

The RHCOS artifacts might not change with every release of OpenShift Container Platform. You must download images with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Only use the appropriate **kernel**, **initramfs**, and **rootfs** artifacts described below for this procedure. RHCOS QCOW2 images are not supported for this installation type.

The file names contain the OpenShift Container Platform version number. They resemble the following examples:

- **kernel:** rhcos-<version>-live-kernel-<architecture>
- **initramfs:** rhcos-<version>-live-initramfs.<architecture>.img
- **rootfs:** rhcos-<version>-live-rootfs.<architecture>.img

4. Upload the **rootfs**, **kernel**, and **initramfs** files to your HTTP server.



IMPORTANT

If you plan to add more compute machines to your cluster after you finish installation, do not delete these files.

5. Configure the network boot infrastructure so that the machines boot from their local disks after RHCOS is installed on them.
6. Configure PXE installation for the RHCOS images and begin the installation. Modify the following example menu entry for your environment and verify that the image and Ignition files are properly accessible:

```
DEFAULT pxeboot
TIMEOUT 20
PROMPT 0
LABEL pxeboot
  KERNEL http://<HTTP_server>/rhcos-<version>-live-kernel-<architecture> 1
  APPEND initrd=http://<HTTP_server>/rhcos-<version>-live-initramfs.<architecture>.img
  coreos.live.rootfs_url=http://<HTTP_server>/rhcos-<version>-live-rootfs.<architecture>.img
  coreos.inst.install_dev=/dev/sda coreos.inst.ignition_url=http://<HTTP_server>/bootstrap.ign
2 3
```

- 1** Specify the location of the live **kernel** file that you uploaded to your HTTP server. The URL must be HTTP, TFTP, or FTP; HTTPS and NFS are not supported.
- 2** If you use multiple NICs, specify a single interface in the **ip** option. For example, to use DHCP on a NIC that is named **eno1**, set **ip=eno1:dhcp**.

3

Specify the locations of the RHCOS files that you uploaded to your HTTP server. The **initrd** parameter value is the location of the **initramfs** file, the **coreos.live.rootfs_url**



NOTE

This configuration does not enable serial console access on machines with a graphical console. To configure a different console, add one or more **console=** arguments to the **APPEND** line. For example, add **console=tty0 console=ttyS0** to set the first PC serial port as the primary console and the graphical console as a secondary console. For more information, see [How does one set up a serial terminal and/or console in Red Hat Enterprise Linux?](#) and "Enabling the serial console for PXE and ISO installation" in the "Advanced RHCOS installation configuration" section.

7. Monitor the progress of the RHCOS installation on the console of the machine.



IMPORTANT


Be sure that the installation is successful on each node before commencing with the OpenShift Container Platform installation. Observing the installation process can also help to determine the cause of RHCOS installation issues that might arise.

8. After RHCOS installs, the system reboots. During reboot, the system applies the Ignition config file that you specified.
9. Check the console output to verify that Ignition ran.

Example command

```
Ignition: ran on 2022/03/14 14:48:33 UTC (this boot)
Ignition: user-provided config was applied
```

10. Continue to create the machines for your cluster.



IMPORTANT

You must create the bootstrap and control plane machines at this time. If the control plane machines are not made schedulable, also create at least two compute machines before you install the cluster.

If the required network, DNS, and load balancer infrastructure are in place, the OpenShift Container Platform bootstrap process begins automatically after the RHCOS nodes have rebooted.



NOTE

RHCOS nodes do not include a default password for the **core** user. You can access the nodes by running **ssh core@<node>.<cluster_name>.<base_domain>** as a user with access to the SSH private key that is paired to the public key that you specified in your **install_config.yaml** file. OpenShift Container Platform 4 cluster nodes running RHCOS are immutable and rely on Operators to apply cluster changes. Accessing cluster nodes by using SSH is not recommended. However, when investigating installation issues, if the OpenShift Container Platform API is not available, or the kubelet is not properly functioning on a target node, SSH access might be required for debugging or disaster recovery.

19.3.11.3. Enabling multipathing with kernel arguments on RHCOS

In OpenShift Container Platform version 4.16, during installation, you can enable multipathing for provisioned nodes. RHCOS supports multipathing on the primary disk. Multipathing provides added benefits of stronger resilience to hardware failure to achieve higher host availability.

During the initial cluster creation, you might want to add kernel arguments to all master or worker nodes. To add kernel arguments to master or worker nodes, you can create a **MachineConfig** object and inject that object into the set of manifest files used by Ignition during cluster setup.

Procedure

1. Change to the directory that contains the installation program and generate the Kubernetes manifests for the cluster:

```
$ ./openshift-install create manifests --dir <installation_directory>
```

2. Decide if you want to add kernel arguments to worker or control plane nodes.
 - Create a machine config file. For example, create a **99-master-kargs-mpath.yaml** that instructs the cluster to add the **master** label and identify the multipath kernel argument:

```
apiVersion: machineconfiguration.openshift.io/v1
kind: MachineConfig
metadata:
  labels:
    machineconfiguration.openshift.io/role: "master"
  name: 99-master-kargs-mpath
spec:
  kernelArguments:
    - 'rd.multipath=default'
    - 'root=/dev/disk/by-label/dm-mpath-root'
```

3. To enable multipathing on worker nodes:
 - Create a machine config file. For example, create a **99-worker-kargs-mpath.yaml** that instructs the cluster to add the **worker** label and identify the multipath kernel argument:


```
apiVersion: machineconfiguration.openshift.io/v1
kind: MachineConfig
metadata:
  labels:
```

```

machineconfiguration.openshift.io/role: "worker"
name: 99-worker-kargs-mpath
spec:
  kernelArguments:
    - 'rd.multipath=default'
    - 'root=/dev/disk/by-label/dm-mpath-root'

```

You can now continue on to create the cluster.



IMPORTANT

Additional postinstallation steps are required to fully enable multipathing. For more information, see “Enabling multipathing with kernel arguments on RHCOS” in *Postinstallation machine configuration tasks*.

In case of MPIO failure, use the `bootlist` command to update the boot device list with alternate logical device names. The command displays a boot list and it designates the possible boot devices for when the system is booted in normal mode.

- a. To display a boot list and specify the possible boot devices if the system is booted in normal mode, enter the following command:

```

$ bootlist -m normal -o
sda

```

- b. To update the boot list for normal mode and add alternate device names, enter the following command:

```

$ bootlist -m normal -o /dev/sdc /dev/sdd /dev/sde
sdc
sdd
sde

```

If the original boot disk path is down, the node reboots from the alternate device registered in the normal boot device list.

19.3.12. Waiting for the bootstrap process to complete

The OpenShift Container Platform bootstrap process begins after the cluster nodes first boot into the persistent RHCOS environment that has been installed to disk. The configuration information provided through the Ignition config files is used to initialize the bootstrap process and install OpenShift Container Platform on the machines. You must wait for the bootstrap process to complete.

Prerequisites

- You have created the Ignition config files for your cluster.
- You have configured suitable network, DNS and load balancing infrastructure.
- You have obtained the installation program and generated the Ignition config files for your cluster.
- You installed RHCOS on your cluster machines and provided the Ignition config files that the OpenShift Container Platform installation program generated.

Procedure

1. Monitor the bootstrap process:

```
$ ./openshift-install --dir <installation_directory> wait-for bootstrap-complete \ 1
--log-level=info 2
```


- 1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.
- 2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Example output

```
INFO Waiting up to 30m0s for the Kubernetes API at https://api.test.example.com:6443...
INFO API v1.29.4 up
INFO Waiting up to 30m0s for bootstrapping to complete...
INFO It is now safe to remove the bootstrap resources
```

The command succeeds when the Kubernetes API server signals that it has been bootstrapped on the control plane machines.

2. After the bootstrap process is complete, remove the bootstrap machine from the load balancer.



IMPORTANT

You must remove the bootstrap machine from the load balancer at this point. You can also remove or reformat the bootstrap machine itself.

19.3.13. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

19.3.14. Approving the certificate signing requests for your machines

When you add machines to a cluster, two pending certificate signing requests (CSRs) are generated for each machine that you added. You must confirm that these CSRs are approved or, if necessary, approve them yourself. The client requests must be approved first, followed by the server requests.

Prerequisites

- You added machines to your cluster.

Procedure

1. Confirm that the cluster recognizes the machines:

```
$ oc get nodes
```

Example output

```
NAME      STATUS   ROLES    AGE   VERSION
master-0  Ready    master   63m   v1.29.4
master-1  Ready    master   63m   v1.29.4
master-2  Ready    master   64m   v1.29.4
```

The output lists all of the machines that you created.



NOTE

The preceding output might not include the compute nodes, also known as worker nodes, until some CSRs are approved.

2. Review the pending CSRs and ensure that you see the client requests with the **Pending** or **Approved** status for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

```
NAME      AGE   REQUESTOR                                     CONDITION
csr-8b2br  15m   system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper  Pending
csr-8vnps  15m   system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper  Pending
...
```

In this example, two machines are joining the cluster. You might see more approved CSRs in the list.

- If the CSRs were not approved, after all of the pending CSRs for the machines you added are in **Pending** status, approve the CSRs for your cluster machines:



NOTE

Because the CSRs rotate automatically, approve your CSRs within an hour of adding the machines to the cluster. If you do not approve them within an hour, the certificates will rotate, and more than two certificates will be present for each node. You must approve all of these certificates. After the client CSR is approved, the Kubelet creates a secondary CSR for the serving certificate, which requires manual approval. Then, subsequent serving certificate renewal requests are automatically approved by the **machine-approver** if the Kubelet requests a new certificate with identical parameters.



NOTE

For clusters running on platforms that are not machine API enabled, such as bare metal and other user-provisioned infrastructure, you must implement a method of automatically approving the kubelet serving certificate requests (CSRs). If a request is not approved, then the **oc exec**, **oc rsh**, and **oc logs** commands cannot succeed, because a serving certificate is required when the API server connects to the kubelet. Any operation that contacts the Kubelet endpoint requires this certificate approval to be in place. The method must watch for new CSRs, confirm that the CSR was submitted by the **node-bootstrap** service account in the **system:node** or **system:admin** groups, and confirm the identity of the node.

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{"\n"}\n{{end}}{{end}}' | xargs --no-run-if-empty oc adm certificate approve
```



NOTE

Some Operators might not become available until some CSRs are approved.

- Now that your client requests are approved, you must review the server requests for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

NAME	AGE	REQUESTOR	CONDITION
csr-bfd72	5m26s	system:node:ip-10-0-50-126.us-east-2.compute.internal	Pending
csr-c57lv	5m26s	system:node:ip-10-0-95-157.us-east-2.compute.internal	Pending
...			

5. If the remaining CSRs are not approved, and are in the **Pending** status, approve the CSRs for your cluster machines:

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

1 **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:


```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{"\n"}{{end}}{{end}}' | xargs oc adm certificate approve
```

6. After all client and server CSRs have been approved, the machines have the **Ready** status. Verify this by running the following command:

```
$ oc get nodes
```

Example output

NAME	STATUS	ROLES	AGE	VERSION
master-0	Ready	master	73m	v1.29.4
master-1	Ready	master	73m	v1.29.4
master-2	Ready	master	74m	v1.29.4
worker-0	Ready	worker	11m	v1.29.4
worker-1	Ready	worker	11m	v1.29.4



NOTE

It can take a few minutes after approval of the server CSRs for the machines to transition to the **Ready** status.

Additional information

- For more information on CSRs, see [Certificate Signing Requests](#).

19.3.15. Initial Operator configuration

After the control plane initializes, you must immediately configure some Operators so that they all become available.

Prerequisites

- Your control plane has initialized.

Procedure

1. Watch the cluster components come online:

```
$ watch -n5 oc get clusteroperators
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.16.0	True	False	False	19m
baremetal	4.16.0	True	False	False	37m
cloud-credential	4.16.0	True	False	False	40m
cluster-autoscaler	4.16.0	True	False	False	37m
config-operator	4.16.0	True	False	False	38m
console	4.16.0	True	False	False	26m
csi-snapshot-controller	4.16.0	True	False	False	37m
dns	4.16.0	True	False	False	37m
etcd	4.16.0	True	False	False	36m
image-registry	4.16.0	True	False	False	31m
ingress	4.16.0	True	False	False	30m
insights	4.16.0	True	False	False	31m
kube-apiserver	4.16.0	True	False	False	26m
kube-controller-manager	4.16.0	True	False	False	36m
kube-scheduler	4.16.0	True	False	False	36m
kube-storage-version-migrator	4.16.0	True	False	False	37m
machine-api	4.16.0	True	False	False	29m
machine-approver	4.16.0	True	False	False	37m
machine-config	4.16.0	True	False	False	36m
marketplace	4.16.0	True	False	False	37m
monitoring	4.16.0	True	False	False	29m
network	4.16.0	True	False	False	38m
node-tuning	4.16.0	True	False	False	37m
openshift-apiserver	4.16.0	True	False	False	32m
openshift-controller-manager	4.16.0	True	False	False	30m
openshift-samples	4.16.0	True	False	False	32m
operator-lifecycle-manager	4.16.0	True	False	False	37m
operator-lifecycle-manager-catalog	4.16.0	True	False	False	37m
operator-lifecycle-manager-packageserver	4.16.0	True	False	False	32m
service-ca	4.16.0	True	False	False	38m
storage	4.16.0	True	False	False	37m

2. Configure the Operators that are not available.

19.3.15.1. Disabling the default OperatorHub catalog sources

Operator catalogs that source content provided by Red Hat and community projects are configured for OperatorHub by default during an OpenShift Container Platform installation. In a restricted network environment, you must disable the default catalogs as a cluster administrator.

Procedure

- Disable the sources for the default catalogs by adding **disableAllDefaultSources: true** to the **OperatorHub** object:

```
$ oc patch OperatorHub cluster --type json \
  -p '[{"op": "add", "path": "/spec/disableAllDefaultSources", "value": true}]'
```

TIP

Alternatively, you can use the web console to manage catalog sources. From the **Administration** → **Cluster Settings** → **Configuration** → **OperatorHub** page, click the **Sources** tab, where you can create, update, delete, disable, and enable individual sources.

19.3.15.2. Image registry storage configuration

The Image Registry Operator is not initially available for platforms that do not provide default storage. After installation, you must configure your registry to use storage so that the Registry Operator is made available.

Instructions are shown for configuring a persistent volume, which is required for production clusters. Where applicable, instructions are shown for configuring an empty directory as the storage location, which is available for only non-production clusters.

Additional instructions are provided for allowing the image registry to use block storage types by using the **Recreate** rollout strategy during upgrades.

19.3.15.2.1. Changing the image registry's management state

To start the image registry, you must change the Image Registry Operator configuration's **managementState** from **Removed** to **Managed**.

Procedure

- Change **managementState** Image Registry Operator configuration from **Removed** to **Managed**. For example:


```
$ oc patch configs.imageregistry.operator.openshift.io cluster --type merge --patch '{"spec": {"managementState": "Managed"}}'
```

19.3.15.2.2. Configuring registry storage for IBM Power

As a cluster administrator, following installation you must configure your registry to use storage.

Prerequisites

- You have access to the cluster as a user with the **cluster-admin** role.
- You have a cluster on IBM Power®.
- You have provisioned persistent storage for your cluster, such as Red Hat OpenShift Data Foundation.



IMPORTANT

OpenShift Container Platform supports **ReadWriteOnce** access for image registry storage when you have only one replica. **ReadWriteOnce** access also requires that the registry uses the **Recreate** rollout strategy. To deploy an image registry that supports high availability with two or more replicas, **ReadWriteMany** access is required.

- Must have 100Gi capacity.

Procedure

1. To configure your registry to use storage, change the **spec.storage.pvc** in the **configs.imageregistry/cluster** resource.



NOTE

When you use shared storage, review your security settings to prevent outside access.

2. Verify that you do not have a registry pod:

```
$ oc get pod -n openshift-image-registry -l docker-registry=default
```

Example output

```
No resources found in openshift-image-registry namespace
```



NOTE

If you do have a registry pod in your output, you do not need to continue with this procedure.

3. Check the registry configuration:

```
$ oc edit configs.imageregistry.operator.openshift.io
```

Example output

```
storage:
  pvc:
    claim:
```

Leave the **claim** field blank to allow the automatic creation of an **image-registry-storage** PVC.

4. Check the **clusteroperator** status:

```
$ oc get clusteroperator image-registry
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
image-registry	4.16	True	False	False	6h50m

5. Ensure that your registry is set to managed to enable building and pushing of images.

- Run:

```
$ oc edit configs.imageregistry/cluster
```

Then, change the line

```
managementState: Removed
```

to

```
managementState: Managed
```

19.3.15.2.3. Configuring storage for the image registry in non-production clusters

You must configure storage for the Image Registry Operator. For non-production clusters, you can set the image registry to an empty directory. If you do so, all images are lost if you restart the registry.

Procedure

- To set the image registry storage to an empty directory:

```
$ oc patch configs.imageregistry.operator.openshift.io cluster --type merge --patch '{"spec": {"storage":{"emptyDir":{}}}'
```



WARNING

Configure this option for only non-production clusters.

If you run this command before the Image Registry Operator initializes its components, the **oc patch** command fails with the following error:

```
Error from server (NotFound): configs.imageregistry.operator.openshift.io "cluster" not found
```

Wait a few minutes and run the command again.

19.3.16. Completing installation on user-provisioned infrastructure

After you complete the Operator configuration, you can finish installing the cluster on infrastructure that you provide.

Prerequisites

- Your control plane has initialized.
- You have completed the initial Operator configuration.

Procedure

1. Confirm that all the cluster components are online with the following command:

```
$ watch -n5 oc get clusteroperators
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.16.0	True	False	False	19m
baremetal	4.16.0	True	False	False	37m
cloud-credential	4.16.0	True	False	False	40m
cluster-autoscaler	4.16.0	True	False	False	37m
config-operator	4.16.0	True	False	False	38m
console	4.16.0	True	False	False	26m
csi-snapshot-controller	4.16.0	True	False	False	37m
dns	4.16.0	True	False	False	37m
etcd	4.16.0	True	False	False	36m
image-registry	4.16.0	True	False	False	31m
ingress	4.16.0	True	False	False	30m
insights	4.16.0	True	False	False	31m
kube-apiserver	4.16.0	True	False	False	26m
kube-controller-manager	4.16.0	True	False	False	36m
kube-scheduler	4.16.0	True	False	False	36m
kube-storage-version-migrator	4.16.0	True	False	False	37m
machine-api	4.16.0	True	False	False	29m
machine-approver	4.16.0	True	False	False	37m
machine-config	4.16.0	True	False	False	36m
marketplace	4.16.0	True	False	False	37m
monitoring	4.16.0	True	False	False	29m
network	4.16.0	True	False	False	38m
node-tuning	4.16.0	True	False	False	37m
openshift-apiserver	4.16.0	True	False	False	32m
openshift-controller-manager	4.16.0	True	False	False	30m
openshift-samples	4.16.0	True	False	False	32m
operator-lifecycle-manager	4.16.0	True	False	False	37m
operator-lifecycle-manager-catalog	4.16.0	True	False	False	37m
operator-lifecycle-manager-packageserver	4.16.0	True	False	False	32m
service-ca	4.16.0	True	False	False	38m
storage	4.16.0	True	False	False	37m

Alternatively, the following command notifies you when all of the clusters are available. It also retrieves and displays credentials:

```
$ ./openshift-install --dir <installation_directory> wait-for install-complete 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

Example output

```
INFO Waiting up to 30m0s for the cluster to initialize...
```

The command succeeds when the Cluster Version Operator finishes deploying the OpenShift Container Platform cluster from Kubernetes API server.



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

2. Confirm that the Kubernetes API server is communicating with the pods.

a. To view a list of all pods, use the following command:

```
$ oc get pods --all-namespaces
```

Example output

```

NAMESPACE           NAME                                     READY  STATUS
RESTARTS  AGE
openshift-apiserver-operator  openshift-apiserver-operator-85cb746d55-zqhs8  1/1
Running    1    9m
openshift-apiserver          apiserver-67b9g                                1/1  Running  0
3m
openshift-apiserver          apiserver-ljcmx                                1/1  Running  0
1m
openshift-apiserver          apiserver-z25h4                                1/1  Running  0
2m
openshift-authentication-operator  authentication-operator-69d5d8bf84-vh2n8      1/1
Running    0    5m
...

```

b. View the logs for a pod that is listed in the output of the previous command by using the following command:

```
$ oc logs <pod_name> -n <namespace> 1
```

- 1** Specify the pod name and namespace, as shown in the output of the previous command.

If the pod logs display, the Kubernetes API server can communicate with the cluster machines.

3. Additional steps are required to enable multipathing. Do not enable multipathing during installation.
See "Enabling multipathing with kernel arguments on RHCOS" in the *Postinstallation machine configuration tasks* documentation for more information.
4. Register your cluster on the [Cluster registration](#) page.

19.3.17. Next steps

- [Enabling multipathing with kernel arguments on RHCOS](#) .
- [Customize your cluster](#) .
- If the mirror registry that you used to install your cluster has a trusted CA, add it to the cluster by [configuring additional trust stores](#) .
- If necessary, you can [opt out of remote health reporting](#) .
- If necessary, see [Registering your disconnected cluster](#)

19.4. INSTALLATION CONFIGURATION PARAMETERS FOR IBM POWER

Before you deploy an OpenShift Container Platform cluster, you provide a customized **install-config.yaml** installation configuration file that describes the details for your environment.

19.4.1. Available installation configuration parameters for IBM Power

The following tables specify the required, optional, and IBM Power-specific installation configuration parameters that you can set as part of the installation process.



NOTE

After installation, you cannot modify these parameters in the **install-config.yaml** file.

19.4.1.1. Required configuration parameters

Required installation configuration parameters are described in the following table:

Table 19.39. Required parameters

Parameter	Description	Values
apiVersion:	The API version for the install-config.yaml content. The current version is v1 . The installation program may also support older API versions.	String

Parameter	Description	Values
baseDomain:	The base domain of your cloud provider. The base domain is used to create routes to your OpenShift Container Platform cluster components. The full DNS name for your cluster is a combination of the baseDomain and metadata.name parameter values that uses the <metadata.name>.<baseDomain> format.	A fully-qualified domain or subdomain name, such as example.com .
metadata:	Kubernetes resource ObjectMeta , from which only the name parameter is consumed.	Object
metadata: name:	The name of the cluster. DNS records for the cluster are all subdomains of {{.metadata.name}} . {{.baseDomain}} .	String of lowercase letters, hyphens (-), and periods (.), such as dev .
platform:	The configuration for the specific platform upon which to perform the installation: aws, baremetal, azure, gcp, ibmcloud, nutanix, openstack, powervs, vsphere , or {} . For additional information about platform . <platform> parameters, consult the table for your specific platform that follows.	Object

Parameter	Description	Values
<code>pullSecret:</code>	Get a pull secret from Red Hat OpenShift Cluster Manager to authenticate downloading container images for OpenShift Container Platform components from services such as Quay.io.	<pre>{ "auths":{ "cloud.openshift.com":{ "auth":"b3Blb=", "email":"you@example.com" }, "quay.io":{ "auth":"b3Blb=", "email":"you@example.com" } } }</pre>

19.4.1.2. Network configuration parameters

You can customize your installation configuration based on the requirements of your existing network infrastructure. For example, you can expand the IP address block for the cluster network or provide different IP address blocks than the defaults.

- If you use the Red Hat OpenShift Networking OVN-Kubernetes network plugin, both IPv4 and IPv6 address families are supported.

If you configure your cluster to use both IP address families, review the following requirements:

- Both IP families must use the same network interface for the default gateway.
- Both IP families must have the default gateway.
- You must specify IPv4 and IPv6 addresses in the same order for all network configuration parameters. For example, in the following configuration IPv4 addresses are listed before IPv6 addresses.


```
networking:
  clusterNetwork:
    - cidr: 10.128.0.0/14
      hostPrefix: 23
    - cidr: fd00:10:128::/56
      hostPrefix: 64
  serviceNetwork:
    - 172.30.0.0/16
    - fd00:172:16::/112
```




NOTE

Globalnet is not supported with Red Hat OpenShift Data Foundation disaster recovery solutions. For regional disaster recovery scenarios, ensure that you use a nonoverlapping range of private IP addresses for the cluster and service networks in each cluster.

Table 19.40. Network parameters

Parameter	Description	Values
<code>networking:</code>	The configuration for the cluster network.	Object  NOTE You cannot modify parameters specified by the networking object after installation.
<code>networking: networkType:</code>	The Red Hat OpenShift Networking network plugin to install.	OVNKubernetes. OVNKubernetes is a CNI plugin for Linux networks and hybrid networks that contain both Linux and Windows servers. The default value is OVNKubernetes .
<code>networking: clusterNetwork:</code>	The IP address blocks for pods. The default value is 10.128.0.0/14 with a host prefix of /23 . If you specify multiple IP address blocks, the blocks must not overlap.	An array of objects. For example: <code>networking: clusterNetwork: - cidr: 10.128.0.0/14 hostPrefix: 23</code>
<code>networking: clusterNetwork: cidr:</code>	Required if you use networking.clusterNetwork . An IP address block. An IPv4 network.	An IP address block in Classless Inter-Domain Routing (CIDR) notation. The prefix length for an IPv4 block is between 0 and 32 .
<code>networking: clusterNetwork: hostPrefix:</code>	The subnet prefix length to assign to each individual node. For example, if hostPrefix is set to 23 then each node is assigned a /23 subnet out of the given cidr . A hostPrefix value of 23 provides 510 ($2^{(32 - 23)} - 2$) pod IP addresses.	A subnet prefix. The default value is 23 .
<code>networking: serviceNetwork:</code>	The IP address block for services. The default value is 172.30.0.0/16 . The OVN-Kubernetes network plugins supports only a single IP address block for the service network.	An array with an IP address block in CIDR format. For example: <code>networking: serviceNetwork: - 172.30.0.0/16</code>

Parameter	Description	Values
<code>networking: machineNetwork:</code>	<p>The IP address blocks for machines.</p> <p>If you specify multiple IP address blocks, the blocks must not overlap.</p> <p>If you specify multiple IP kernel arguments, the machineNetwork.cidr value must be the CIDR of the primary network.</p>	<p>An array of objects. For example:</p> <pre>networking: machineNetwork: - cidr: 10.0.0.0/16</pre>
<code>networking: machineNetwork: cidr:</code>	<p>Required if you use networking.machineNetwork. An IP address block. The default value is 10.0.0.0/16 for all platforms other than libvirt and IBM Power® Virtual Server. For libvirt, the default value is 192.168.126.0/24. For IBM Power® Virtual Server, the default value is 192.168.0.0/24.</p>	<p>An IP network block in CIDR notation.</p> <p>For example, 10.0.0.0/16.</p> <div style="display: flex; align-items: center;">  <div> <p>NOTE</p> <p>Set the networking.machineNetwork to match the CIDR that the preferred NIC resides in.</p> </div> </div>


19.4.1.3. Optional configuration parameters


Optional installation configuration parameters are described in the following table:

Table 19.41. Optional parameters

Parameter	Description	Values
<code>additionalTrustBundle:</code>	<p>A PEM-encoded X.509 certificate bundle that is added to the nodes' trusted certificate store. This trust bundle may also be used when a proxy has been configured.</p>	String
<code>capabilities:</code>	<p>Controls the installation of optional core cluster components. You can reduce the footprint of your OpenShift Container Platform cluster by disabling optional components. For more information, see the "Cluster capabilities" page in <i>Installing</i>.</p>	String array


Parameter	Description	Values
capabilities: baselineCapabilitySet:	Selects an initial set of optional capabilities to enable. Valid values are None , v4.11 , v4.12 and vCurrent . The default value is vCurrent .	String
capabilities: additionalEnabledCapabilities:	Extends the set of optional capabilities beyond what you specify in baselineCapabilitySet . You may specify multiple capabilities in this parameter.	String array
cpuPartitioningMode:	Enables workload partitioning, which isolates OpenShift Container Platform services, cluster management workloads, and infrastructure pods to run on a reserved set of CPUs. Workload partitioning can only be enabled during installation and cannot be disabled after installation. While this field enables workload partitioning, it does not configure workloads to use specific CPUs. For more information, see the <i>Workload partitioning</i> page in the <i>Scalability and Performance</i> section.	None or AllNodes . None is the default value.
compute:	The configuration for the machines that comprise the compute nodes.	Array of MachinePool objects.
compute: architecture:	Determines the instruction set architecture of the machines in the pool. Currently, heterogeneous clusters are not supported, so all pools must specify the same architecture. Valid values are ppc64le (the default).	String

Parameter	Description	Values
compute: hyperthreading:	<p>Whether to enable or disable simultaneous multithreading, or hyperthreading, on compute machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.</p> <div style="display: flex; align-items: center;">  <div> <p>IMPORTANT</p> <p>If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.</p> </div> </div>	Enabled or Disabled
compute: name:	Required if you use compute . The name of the machine pool.	worker
compute: platform:	Required if you use compute . Use this parameter to specify the cloud provider to host the worker machines. This parameter value must match the controlPlane.platform parameter value.	aws, azure, gcp, ibmcloud, nutanix, openstack, powervs, vsphere , or {}
compute: replicas:	The number of compute machines, which are also known as worker machines, to provision.	A positive integer greater than or equal to 2 . The default value is 3 .
featureSet:	Enables the cluster for a feature set. A feature set is a collection of OpenShift Container Platform features that are not enabled by default. For more information about enabling a feature set during installation, see "Enabling features using feature gates".	String. The name of the feature set to enable, such as TechPreviewNoUpgrade .
controlPlane:	The configuration for the machines that comprise the control plane.	Array of MachinePool objects.

Parameter	Description	Values
<code>controlPlane: architecture:</code>	Determines the instruction set architecture of the machines in the pool. Currently, heterogeneous clusters are not supported, so all pools must specify the same architecture. Valid values are ppc64le (the default).	String
<code>controlPlane: hyperthreading:</code>	Whether to enable or disable simultaneous multithreading, or hyperthreading , on control plane machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.  IMPORTANT If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.	Enabled or Disabled
<code>controlPlane: name:</code>	Required if you use controlPlane . The name of the machine pool.	master
<code>controlPlane: platform:</code>	Required if you use controlPlane . Use this parameter to specify the cloud provider that hosts the control plane machines. This parameter value must match the compute.platform parameter value.	aws, azure, gcp, ibmcloud, nutanix, openstack, powervs, vsphere , or <code>{}</code>
<code>controlPlane: replicas:</code>	The number of control plane machines to provision.	Supported values are 3 , or 1 when deploying single-node OpenShift.
<code>credentialsMode:</code>	The Cloud Credential Operator (CCO) mode. If no mode is specified, the CCO dynamically tries to determine the capabilities of the provided credentials, with a preference for mint mode on the platforms where multiple modes are supported.	Mint, Passthrough, Manual or an empty string (<code>""</code>). ^[1]

Parameter	Description	Values
	<p>Enable or disable FIPS mode. The default is <code>false</code> (disabled). If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.</p> <div data-bbox="488 479 592 1527" style="background-color: black; color: white; padding: 5px; margin-bottom: 10px;"> <p>IMPORTANT</p> <p>To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see Installing the system in FIPS mode.</p> <p>When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.</p> </div> <div data-bbox="488 1576 592 1771" style="background-color: black; color: white; padding: 5px;"> <p>NOTE</p> <p>If you are using Azure File storage, you cannot enable FIPS mode.</p> </div>	<p><code>false</code> or <code>true</code></p>

Parameter	Description	Values
<code>imageContentSources:</code>	Sources and repositories for the release-image content.	Array of objects. Includes a source and, optionally, mirrors , as described in the following rows of this table.
<code>imageContentSources: source:</code>	Required if you use imageContentSources . Specify the repository that users refer to, for example, in image pull specifications.	String
<code>imageContentSources: mirrors:</code>	Specify one or more repositories that may also contain the same images.	Array of strings
<code>publish:</code>	How to publish or expose the user-facing endpoints of your cluster, such as the Kubernetes API, OpenShift routes.	<p>Internal or External. The default value is External.</p> <p>Setting this field to Internal is not supported on non-cloud platforms.</p> <div data-bbox="975 1451 1082 1706" style="display: inline-block; vertical-align: middle;">  </div> <p>IMPORTANT</p> <p>If the value of the field is set to Internal, the cluster will become non-functional. For more information, refer to BZ#1953035.</p>

Parameter	Description	Values
sshKey:	<p>The SSH key to authenticate access to your cluster machines.</p>  <p>NOTE</p> <p>For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your ssh-agent process uses.</p>	<p>For example, sshKey: ssh-ed25519 AAAA...</p>

1. Not all CCO modes are supported for all cloud providers. For more information about CCO modes, see the "Managing cloud provider credentials" entry in the *Authentication and authorization* content.

CHAPTER 20. INSTALLING ON IBM POWER VIRTUAL SERVER

20.1. PREPARING TO INSTALL ON IBM POWER VIRTUAL SERVER

The installation workflows documented in this section are for IBM Power® Virtual Server infrastructure environments.

20.1.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).

20.1.2. Requirements for installing OpenShift Container Platform on IBM Power Virtual Server

Before installing OpenShift Container Platform on IBM Power® Virtual Server you must create a service account and configure an IBM Cloud® account. See [Configuring an IBM Cloud® account](#) for details about creating an account, configuring DNS and supported IBM Power® Virtual Server regions.

You must manually manage your cloud credentials when installing a cluster to IBM Power® Virtual Server. Do this by configuring the Cloud Credential Operator (CCO) for manual mode before you install the cluster.

20.1.3. Choosing a method to install OpenShift Container Platform on IBM Power Virtual Server

You can install OpenShift Container Platform on IBM Power® Virtual Server using installer-provisioned infrastructure. This process involves using an installation program to provision the underlying infrastructure for your cluster. Installing OpenShift Container Platform on IBM Power® Virtual Server using user-provisioned infrastructure is not supported at this time.

See [Installation process](#) for more information about installer-provisioned installation processes.

20.1.3.1. Installing a cluster on installer-provisioned infrastructure

You can install a cluster on IBM Power® Virtual Server infrastructure that is provisioned by the OpenShift Container Platform installation program by using one of the following methods:


- **Installing a customized cluster on IBM Power® Virtual Server** You can install a customized cluster on IBM Power® Virtual Server infrastructure that the installation program provisions. The installation program allows for some customization to be applied at the installation stage. Many other customization options are available [post-installation](#).
- **Installing a cluster on IBM Power® Virtual Server into an existing VPC** You can install OpenShift Container Platform on IBM Power® Virtual Server into an existing Virtual Private Cloud (VPC). You can use this installation method if you have constraints set by the guidelines of your company, such as limits when creating new accounts or infrastructure.
- **Installing a private cluster on IBM Power® Virtual Server** You can install a private cluster on IBM Power® Virtual Server. You can use this method to deploy OpenShift Container Platform on an internal network that is not visible to the internet.

- **Installing a cluster on IBM Power® Virtual Server in a restricted network** You can install OpenShift Container Platform on IBM Power® Virtual Server on installer-provisioned infrastructure by using an internal mirror of the installation release content. You can use this method to install a cluster that does not require an active internet connection to obtain the software components.

20.1.4. Configuring the Cloud Credential Operator utility

The Cloud Credential Operator (CCO) manages cloud provider credentials as Kubernetes custom resource definitions (CRDs). To install a cluster on IBM Power® Virtual Server, you must set the CCO to **manual** mode as part of the installation process.

To create and manage cloud credentials from outside of the cluster when the Cloud Credential Operator (CCO) is operating in manual mode, extract and prepare the CCO utility (**ccoctl**) binary.



NOTE

The **ccoctl** utility is a Linux binary that must run in a Linux environment.

Prerequisites

- You have access to an OpenShift Container Platform account with cluster administrator access.
- You have installed the OpenShift CLI (**oc**).

Procedure

1. Set a variable for the OpenShift Container Platform release image by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk 'release image/ {print $3}')
```

2. Obtain the CCO container image from the OpenShift Container Platform release image by running the following command:

```
$ CCO_IMAGE=$(oc adm release info --image-for='cloud-credential-operator'
$RELEASE_IMAGE -a ~/.pull-secret)
```



NOTE

Ensure that the architecture of the **\$RELEASE_IMAGE** matches the architecture of the environment in which you will use the **ccoctl** tool.

3. Extract the **ccoctl** binary from the CCO container image within the OpenShift Container Platform release image by running the following command:

```
$ oc image extract $CCO_IMAGE \
--file="/usr/bin/ccoctl.<rhel_version>" \ 1
-a ~/.pull-secret
```

- 1** For **<rhel_version>**, specify the value that corresponds to the version of Red Hat Enterprise Linux (RHEL) that the host uses. If no value is specified, **ccoctl.rhel8** is used by default. The following values are valid:

- **rhel8**: Specify this value for hosts that use RHEL 8.
- **rhel9**: Specify this value for hosts that use RHEL 9.

4. Change the permissions to make **ccoctl** executable by running the following command:

```
$ chmod 775 ccoctl.<rhel_version>
```

Verification

- To verify that **ccoctl** is ready to use, display the help file by running the following command:

```
$ ccoctl --help
```

Output of **ccoctl --help**

```
OpenShift credentials provisioning tool
```

```
Usage:
```

```
ccoctl [command]
```

```
Available Commands:
```

```
aws      Manage credentials objects for AWS cloud
azure    Manage credentials objects for Azure
gcp      Manage credentials objects for Google cloud
help     Help about any command
ibmcloud Manage credentials objects for IBM Cloud
nutanix  Manage credentials objects for Nutanix
```

```
Flags:
```

```
-h, --help  help for ccoctl
```

```
Use "ccoctl [command] --help" for more information about a command.
```

Additional resources

- [Rotating API keys](#)

20.1.5. Next steps

- [Configuring an IBM Cloud® account](#)

20.2. CONFIGURING AN IBM CLOUD ACCOUNT

Before you can install OpenShift Container Platform, you must configure an IBM Cloud® account.

20.2.1. Prerequisites

- You have an IBM Cloud® account with a subscription. You cannot install OpenShift Container Platform on a free or on a trial IBM Cloud® account.

20.2.2. Quotas and limits on IBM Power Virtual Server

The OpenShift Container Platform cluster uses several IBM Cloud® and IBM Power® Virtual Server components, and the default quotas and limits affect your ability to install OpenShift Container Platform clusters. If you use certain cluster configurations, deploy your cluster in certain regions, or run multiple clusters from your account, you might need to request additional resources for your IBM Cloud® account.

For a comprehensive list of the default IBM Cloud® quotas and service limits, see the IBM Cloud® documentation for [Quotas and service limits](#).

Virtual Private Cloud

Each OpenShift Container Platform cluster creates its own Virtual Private Cloud (VPC). The default quota of VPCs per region is 10. If you have 10 VPCs created, you will need to increase your quota before attempting an installation.

Application load balancer

By default, each cluster creates two application load balancers (ALBs):

- Internal load balancer for the control plane API server
- External load balancer for the control plane API server

You can create additional **LoadBalancer** service objects to create additional ALBs. The default quota of VPC ALBs are 50 per region. To have more than 50 ALBs, you must increase this quota.

VPC ALBs are supported. Classic ALBs are not supported for IBM Power® Virtual Server.

Transit Gateways

Each OpenShift Container Platform cluster creates its own Transit Gateway to enable communication with a VPC. The default quota of transit gateways per account is 10. If you have 10 transit gateways created, you will need to increase your quota before attempting an installation.

Dynamic Host Configuration Protocol Service

There is a limit of one Dynamic Host Configuration Protocol (DHCP) service per IBM Power® Virtual Server instance.

Networking

Due to networking limitations, there is a restriction of one OpenShift cluster installed through IPI per zone per account. This is not configurable.

Virtual Server Instances

By default, a cluster creates server instances with the following resources :

- 0.5 CPUs
- 32 GB RAM
- System Type: **s922**
- Processor Type: **uncapped, shared**
- Storage Tier: **Tier-3**

The following nodes are created:

- One bootstrap machine, which is removed after the installation is complete

- Three control plane nodes
- Three compute nodes

For more information, see [Creating a Power Systems Virtual Server](#) in the IBM Cloud® documentation.

20.2.3. Configuring DNS resolution

How you configure DNS resolution depends on the type of OpenShift Container Platform cluster you are installing:

- If you are installing a public cluster, you use IBM Cloud® Internet Services (CIS).
- If you are installing a private cluster, you use IBM Cloud® DNS Services (DNS Services).

20.2.4. Using IBM Cloud Internet Services for DNS resolution

The installation program uses IBM Cloud® Internet Services (CIS) to configure cluster DNS resolution and provide name lookup for a public cluster.



NOTE

This offering does not support IPv6, so dual stack or IPv6 environments are not possible.

You must create a domain zone in CIS in the same account as your cluster. You must also ensure the zone is authoritative for the domain. You can do this using a root domain or subdomain.

Prerequisites

- You have installed the [IBM Cloud® CLI](#).
- You have an existing domain and registrar. For more information, see the IBM® [documentation](#).

Procedure

1. Create a CIS instance to use with your cluster:

- a. Install the CIS plugin:

```
$ ibmcloud plugin install cis
```

- b. Log in to IBM Cloud® by using the CLI:

```
$ ibmcloud login
```

- c. Create the CIS instance:

```
$ ibmcloud cis instance-create <instance_name> standard 1
```

- 1** At a minimum, a **Standard** plan is required for CIS to manage the cluster subdomain and its DNS records.

2. Connect an existing domain to your CIS instance:

- a. Set the context instance for CIS:

```
$ ibmcloud cis instance-set <instance_CRN> 1
```

- 1 The instance CRN (Cloud Resource Name). For example: **ibmcloud cis instance-set crn:v1:bluemix:public:power-iaas:osa21:a/65b64c1f1c29460d8c2e4bbfbd893c2c:c09233ac-48a5-4ccb-a051-d1cfb3fc7eb5::**

- b. Add the domain for CIS:

```
$ ibmcloud cis domain-add <domain_name> 1
```

- 1 The fully qualified domain name. You can use either the root domain or subdomain value as the domain name, depending on which you plan to configure.



NOTE

A root domain uses the form **openshiftcorp.com**. A subdomain uses the form **clusters.openshiftcorp.com**.

- Open the [CIS web console](#), navigate to the **Overview** page, and note your CIS name servers. These name servers will be used in the next step.
- Configure the name servers for your domains or subdomains at the domain's registrar or DNS provider. For more information, see the IBM Cloud® [documentation](#).

20.2.5. IBM Cloud IAM Policies and API Key

To install OpenShift Container Platform into your IBM Cloud® account, the installation program requires an IAM API key, which provides authentication and authorization to access IBM Cloud® service APIs. You can use an existing IAM API key that contains the required policies or create a new one.

For an IBM Cloud® IAM overview, see the IBM Cloud® [documentation](#).

20.2.5.1. Pre-requisite permissions

Table 20.1. Pre-requisite permissions

Role	Access
Viewer, Operator, Editor, Administrator, Reader, Writer, Manager	Internet Services service in <resource_group> resource group
Viewer, Operator, Editor, Administrator, User API key creator, Service ID creator	IAM Identity Service service

Role	Access
Viewer, Operator, Administrator, Editor, Reader, Writer, Manager, Console Administrator	VPC Infrastructure Services service in <resource_group> resource group
Viewer	Resource Group: Access to view the resource group itself. The resource type should equal Resource group , with a value of <your_resource_group_name>.

20.2.5.2. Cluster-creation permissions

Table 20.2. Cluster-creation permissions

Role	Access
Viewer	<resource_group> (Resource Group Created for Your Team)
Viewer, Operator, Editor, Reader, Writer, Manager	All Identity and IAM enabled services in Default resource group
Viewer, Reader	Internet Services service
Viewer, Operator, Reader, Writer, Manager, Content Reader, Object Reader, Object Writer, Editor	Cloud Object Storage service
Viewer	Default resource group: The resource type should equal Resource group , with a value of Default . If your account administrator changed your account's default resource group to something other than Default, use that value instead.
Viewer, Operator, Editor, Reader, Manager	Workspace for IBM Power® Virtual Server service in <resource_group> resource group
Viewer, Operator, Editor, Reader, Writer, Manager, Administrator	Internet Services service in <resource_group> resource group: CIS functional scope string equals reliability
Viewer, Operator, Editor	Transit Gateway service
Viewer, Operator, Editor, Administrator, Reader, Writer, Manager, Console Administrator	VPC Infrastructure Services service <resource_group> resource group

20.2.5.3. Access policy assignment

In IBM Cloud® IAM, access policies can be attached to different subjects:

- Access group (Recommended)
- Service ID
- User

The recommended method is to define IAM access policies in an [access group](#). This helps organize all the access required for OpenShift Container Platform and enables you to onboard users and service IDs to this group. You can also assign access to [users and service IDs](#) directly, if desired.

20.2.5.4. Creating an API key

You must create a user API key or a service ID API key for your IBM Cloud® account.

Prerequisites

- You have assigned the required access policies to your IBM Cloud® account.
- You have attached your IAM access policies to an access group, or other appropriate resource.

Procedure

- Create an API key, depending on how you defined your IAM access policies. For example, if you assigned your access policies to a user, you must create a [user API key](#). If you assigned your access policies to a service ID, you must create a [service ID API key](#). If your access policies are assigned to an access group, you can use either API key type. For more information on IBM Cloud® API keys, see [Understanding API keys](#).

20.2.6. Supported IBM Power Virtual Server regions and zones

You can deploy an OpenShift Container Platform cluster to the following regions:

- **dal** (Dallas, USA)
 - **dal10**
 - **dal12**
- **eu-de** (Frankfurt, Germany)
 - **eu-de-1**
 - **eu-de-2**
- **lon** (London, UK)
 - **lon04**
- **mad** (Madrid, Spain)
 - **mad02**
 - **mad04**
- **osa** (Osaka, Japan)
 - **osa21**

- **sao** (Sao Paulo, Brazil)
 - **sao01**
 - **sao04**
- **syd** (Sydney, Australia)
 - **syd04**
- **wdc** (Washington DC, USA)
 - **wdc06**
 - **wdc07**

You might optionally specify the IBM Cloud® region in which the installer will create any VPC components. Supported regions in IBM Cloud® are:

- **us-south**
- **eu-de**
- **eu-es**
- **eu-gb**
- **jp-osa**
- **au-syd**
- **br-sao**
- **ca-tor**
- **jp-tok**

20.2.7. Next steps

- [Creating an IBM Power® Virtual Server workspace](#)

20.3. CREATING AN IBM POWER VIRTUAL SERVER WORKSPACE

20.3.1. Creating an IBM Power Virtual Server workspace

Use the following procedure to create an IBM Power® Virtual Server workspace.

Procedure

1. To create an IBM Power® Virtual Server workspace, complete step 1 to step 5 from the IBM Cloud® documentation for [Creating an IBM Power® Virtual Server](#) .
2. After it has finished provisioning, retrieve the 32-character alphanumeric Globally Unique Identifier (GUID) of your new workspace by entering the following command:

```
$ ibmcloud resource service-instance <workspace name>
```

20.3.2. Next steps

- [Installing a cluster on IBM Power® Virtual Server with customizations](#)

20.4. INSTALLING A CLUSTER ON IBM POWER VIRTUAL SERVER WITH CUSTOMIZATIONS

In OpenShift Container Platform version 4.16, you can install a customized cluster on infrastructure that the installation program provisions on IBM Power Virtual Server. To customize the installation, you modify parameters in the **install-config.yaml** file before you install the cluster.

20.4.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [configured an IBM Cloud® account](#) to host the cluster.
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.
- You configured the **ccoctl** utility before you installed the cluster. For more information, see [Configuring the Cloud Credential Operator utility](#).

20.4.2. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.



IMPORTANT

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

20.4.3. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added

to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.



NOTE

You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.



NOTE

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

20.4.4. Obtaining the installation program

Before you install OpenShift Container Platform, download the installation file on the host you are using for installation.

Prerequisites

- You have a computer that runs Linux or macOS, with 500 MB of local disk space.


Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select your infrastructure provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.



IMPORTANT

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.



IMPORTANT

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.

4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#) . This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

20.4.5. Exporting the API key

You must set the API key you created as a global variable; the installation program ingests the variable during startup to set the API key.

Prerequisites

- You have created either a user API key or service ID API key for your IBM Cloud® account.

Procedure

- Export your API key for your account as a global variable:

```
$ export IBMCLLOUD_API_KEY=<api_key>
```



IMPORTANT

You must set the variable name exactly as specified; the installation program expects the variable name to be present during startup.

20.4.6. Creating the installation configuration file

You can customize the OpenShift Container Platform cluster you install on

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure


1. Create the **install-config.yaml** file.
 - a. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install create install-config --dir <installation_directory> 1
```

- 1 For **<installation_directory>**, specify the directory name to store the files that the installation program creates.

When specifying the directory:

- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
- Use an empty directory. Some installation assets, such as bootstrap X.509 certificates, have short expiration intervals, therefore you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.



NOTE

Always delete the **~/.powersvs** directory to avoid reusing a stale configuration. Run the following command:

```
$ rm -rf ~/.powersvs
```


- b. At the prompts, provide the configuration details for your cloud:
 - i. Optional: Select an SSH key to use to access your cluster machines.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- ii. Select **powersvs** as the platform to target.
 - iii. Select the region to deploy the cluster to.
 - iv. Select the zone to deploy the cluster to.
 - v. Select the base domain to deploy the cluster to. The base domain corresponds to the public DNS zone that you created for your cluster.
 - vi. Enter a descriptive name for your cluster.
2. Modify the **install-config.yaml** file. You can find more information about the available parameters in the "Installation configuration parameters" section.
 3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.

**IMPORTANT**


The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

Additional resources

- [Installation configuration parameters for IBM Power® Virtual Server](#)

20.4.6.1. Sample customized install-config.yaml file for IBM Power Virtual Server

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.

**IMPORTANT**

This sample YAML file is provided for reference only. You must obtain your **install-config.yaml** file by using the installation program and modify it.

```

apiVersion: v1
baseDomain: example.com
compute: ① ②
- architecture: ppc64le
  hyperthreading: Enabled ③
  name: worker
  platform:
    powervs:
      smtLevel: 8 ④
  replicas: 3
controlPlane: ⑤ ⑥
architecture: ppc64le
hyperthreading: Enabled ⑦
name: master
platform:
  powervs:
    smtLevel: 8 ⑧
  replicas: 3
metadata:
  creationTimestamp: null
  name: example-cluster-name
networking:
  clusterNetwork:
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  machineNetwork:
  - cidr: 192.168.0.0/24
  networkType: OVNKubernetes ⑨
  serviceNetwork:
  - 172.30.0.0/16
platform:
  powervs:
    userID: ibm-user-id
    region: powervs-region
    zone: powervs-zone

```



```

powervsResourceGroup: "ibmcloud-resource-group" 10
serviceInstanceGUID: "powervs-region-service-instance-guid"
vpcRegion : vpc-region
publish: External
pullSecret: '{"auths": ...}' 11
sshKey: ssh-ed25519 AAAA... 12

```

- 1 5** If you do not provide these parameters and values, the installation program provides the default value.
- 2 6** The **controlPlane** section is a single mapping, but the compute section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Although both sections currently define a single machine pool, it is possible that OpenShift Container Platform will support defining multiple compute pools during installation. Only one control plane pool is used.
- 3 7** Whether to enable or disable simultaneous multithreading, or **hyperthreading**. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores. You can disable it by setting the parameter value to **Disabled**. If you disable simultaneous multithreading in some cluster machines, you must disable it in all cluster machines.



IMPORTANT

If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.

- 4 8** The `smtLevel` specifies the level of SMT to set to the control plane and compute machines. The supported values are 1, 2, 4, 8, **'off'** and **'on'**. The default value is 8. The `smtLevel 'off'` sets SMT to off and `smtlevel 'on'` sets SMT to the default value 8 on the cluster nodes.



NOTE

When simultaneous multithreading (SMT), or hyperthreading is not enabled, one vCPU is equivalent to one physical core. When enabled, total vCPUs is computed as: (Thread(s) per core * Core(s) per socket) * Socket(s). The `smtLevel` controls the threads per core. Lower SMT levels may require additional assigned cores when deploying the cluster nodes. You can do this by setting the **'processors'** parameter in the **install-config.yaml** file to an appropriate value to meet the requirements for deploying OpenShift Container Platform successfully.

- 9** The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 10** The name of an existing resource group.
- 11** Required. The installation program prompts you for this value.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

20.4.6.2. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then

creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.

- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.



NOTE

The installation program does not support the proxy **readinessEndpoints** field.



NOTE

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.



NOTE

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

20.4.7. Manually creating IAM

Installing the cluster requires that the Cloud Credential Operator (CCO) operate in manual mode. While the installation program configures the CCO for manual mode, you must specify the identity and access management secrets for your cloud provider.

You can use the Cloud Credential Operator (CCO) utility (**ccoctl**) to create the required IBM Cloud® resources.

Prerequisites

- You have configured the **ccoctl** binary.
- You have an existing **install-config.yaml** file.

Procedure

1. Edit the **install-config.yaml** configuration file so that it contains the **credentialsMode** parameter set to **Manual**.

Example `install-config.yaml` configuration file

```
apiVersion: v1
baseDomain: cluster1.example.com
credentialsMode: Manual 1
compute:
  - architecture: ppc64le
    hyperthreading: Enabled
```

- 1** This line is added to set the **credentialsMode** parameter to **Manual**.
- To generate the manifests, run the following command from the directory that contains the installation program:

```
$ openshift-install create manifests --dir <installation_directory>
```

- From the directory that contains the installation program, set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

- Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included 1 \
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml 2 \
  --to=<path_to_directory_for_credentials_requests> 3
```

- 1** The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2** Specify the location of the **install-config.yaml** file.
- 3** Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

This command creates a YAML file for each **CredentialsRequest** object.

Sample **CredentialsRequest** object

```
apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  labels:
    controller-tools.k8s.io: "1.0"
  name: openshift-image-registry-ibmcos
  namespace: openshift-cloud-credential-operator
spec:
  secretRef:
```

```

name: installer-cloud-credentials
namespace: openshift-image-registry
providerSpec:
  apiVersion: cloudcredential.openshift.io/v1
  kind: IBMCloudProviderSpec
  policies:
  - attributes:
    - name: serviceName
      value: cloud-object-storage
    roles:
    - crn:v1:bluemix:public:iam::::role:Viewer
    - crn:v1:bluemix:public:iam::::role:Operator
    - crn:v1:bluemix:public:iam::::role:Editor
    - crn:v1:bluemix:public:iam::::serviceRole:Reader
    - crn:v1:bluemix:public:iam::::serviceRole:Writer
  - attributes:
    - name: resourceType
      value: resource-group
    roles:
    - crn:v1:bluemix:public:iam::::role:Viewer

```

5. Create the service ID for each credential request, assign the policies defined, create an API key, and generate the secret:

```

$ ccoctl ibmcloud create-service-id \
  --credentials-requests-dir=<path_to_credential_requests_directory> \ 1
  --name=<cluster_name> \ 2
  --output-dir=<installation_directory> \ 3
  --resource-group-name=<resource_group_name> \ 4

```

- 1** Specify the directory containing the files for the component **CredentialsRequest** objects.
- 2** Specify the name of the OpenShift Container Platform cluster.
- 3** Optional: Specify the directory in which you want the **ccoctl** utility to create objects. By default, the utility creates objects in the directory in which the commands are run.
- 4** Optional: Specify the name of the resource group used for scoping the access policies.



NOTE

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

If an incorrect resource group name is provided, the installation fails during the bootstrap phase. To find the correct resource group name, run the following command:

```

$ grep resourceGroup <installation_directory>/manifests/cluster-infrastructure-02-config.yml


```

Verification

- Ensure that the appropriate secrets were generated in your cluster's **manifests** directory.

20.4.8. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.

Procedure

- Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

- 1 For **<installation_directory>**, specify the location of your customized **./install-config.yaml** file.
- 2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to **<installation_directory>/openshift_install.log**.



IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
```

```
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

20.4.9. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.



IMPORTANT

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**.
To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.



NOTE

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your **PATH**.
To check your **PATH**, open a terminal and execute the following command:


```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

20.4.10. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

Additional resources

- [Accessing the web console](#)

20.4.11. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- [About remote health monitoring](#)

20.4.12. Next steps

- [Customize your cluster](#)
- If necessary, you can [opt out of remote health reporting](#)

20.5. INSTALLING A CLUSTER ON IBM POWER VIRTUAL SERVER INTO AN EXISTING VPC

In OpenShift Container Platform version 4.16, you can install a cluster into an existing Virtual Private Cloud (VPC) on IBM Cloud®. The installation program provisions the rest of the required infrastructure, which you can then further customize. To customize the installation, you modify parameters in the **install-config.yaml** file before you install the cluster.

20.5.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [configured an IBM Cloud® account](#) to host the cluster.
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.
- You configured the **ccoctl** utility before you installed the cluster. For more information, see [Configuring the Cloud Credential Operator utility](#).

20.5.2. About using a custom VPC

In OpenShift Container Platform 4.16, you can deploy a cluster using an existing IBM® Virtual Private Cloud (VPC).

Because the installation program cannot know what other components are in your existing subnets, it cannot choose subnet CIDRs and so forth. You must configure networking for the subnets to which you will install the cluster.

20.5.2.1. Requirements for using your VPC

You must correctly configure the existing VPC and its subnets before you install the cluster. The installation program does not create a VPC or VPC subnet in this scenario.

The installation program cannot:

- Subdivide network ranges for the cluster to use
- Set route tables for the subnets
- Set VPC options like DHCP

**NOTE**

The installation program requires that you use the cloud-provided DNS server. Using a custom DNS server is not supported and causes the installation to fail.

20.5.2.2. VPC validation

The VPC and all of the subnets must be in an existing resource group. The cluster is deployed to this resource group.

As part of the installation, specify the following in the **install-config.yaml** file:

- The name of the resource group
- The name of VPC
- The name of the VPC subnet

To ensure that the subnets that you provide are suitable, the installation program confirms that all of the subnets you specify exists.

**NOTE**

Subnet IDs are not supported.

20.5.2.3. Isolation between clusters

If you deploy OpenShift Container Platform to an existing network, the isolation of cluster services is reduced in the following ways:

- ICMP Ingress is allowed to the entire network.
- TCP port 22 Ingress (SSH) is allowed to the entire network.
- Control plane TCP 6443 Ingress (Kubernetes API) is allowed to the entire network.
- Control plane TCP 22623 Ingress (MCS) is allowed to the entire network.

20.5.3. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.

**IMPORTANT**

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

20.5.4. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.

**IMPORTANT**

Do not skip this procedure in production environments, where disaster recovery and debugging is required.

**NOTE**

You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.

**NOTE**

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

- View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

- Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.



NOTE

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- If the `ssh-agent` process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```



NOTE

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

- Add your SSH private key to the `ssh-agent`:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

20.5.5. Obtaining the installation program

Before you install OpenShift Container Platform, download the installation file on the host you are using for installation.

Prerequisites

- You have a computer that runs Linux or macOS, with 500 MB of local disk space.

Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select your infrastructure provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.



IMPORTANT

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.



IMPORTANT

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.

4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#) . This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

20.5.6. Exporting the API key

You must set the API key you created as a global variable; the installation program ingests the variable during startup to set the API key.


Prerequisites

- You have created either a user API key or service ID API key for your IBM Cloud® account.

Procedure

- Export your API key for your account as a global variable:

```
$ export IBM_CLOUD_API_KEY=<api_key>
```

**IMPORTANT**

You must set the variable name exactly as specified; the installation program expects the variable name to be present during startup.

20.5.7. Creating the installation configuration file

You can customize the OpenShift Container Platform cluster you install on

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create the **install-config.yaml** file.
 - a. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install create install-config --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the directory name to store the files that the installation program creates.

When specifying the directory:


- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
- Use an empty directory. Some installation assets, such as bootstrap X.509 certificates, have short expiration intervals, therefore you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

- b. At the prompts, provide the configuration details for your cloud:
 - i. Optional: Select an SSH key to use to access your cluster machines.

**NOTE**

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- ii. Enter a descriptive name for your cluster.
2. Modify the **install-config.yaml** file. You can find more information about the available parameters in the "Installation configuration parameters" section.
 3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.

**IMPORTANT**

The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

Additional resources

- [Installation configuration parameters for IBM Power® Virtual Server](#)

20.5.7.1. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 20.3. Minimum resource requirements

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Bootstrap	RHCOS	4	16 GB	100 GB	300
Control plane	RHCOS	4	16 GB	100 GB	300
Compute	RHCOS, RHEL 8.6 and later [3]	2	8 GB	100 GB	300

1. One vCPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio: $(\text{threads per core} \times \text{cores}) \times \text{sockets} = \text{vCPUs}$.
2. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes which require a 10 ms p99 fsync duration. Note that on many cloud platforms, storage size and IOPS scale together, so you might need to over-allocate storage volume to obtain sufficient performance.
3. As with all user-provisioned installations, if you choose to use RHEL compute machines in your cluster, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Use of RHEL 7 compute machines is deprecated and has been removed in OpenShift Container Platform 4.10 and later.

**NOTE**

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

20.5.7.2. Sample customized install-config.yaml file for IBM Power Virtual Server

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.

**IMPORTANT**

This sample YAML file is provided for reference only. You must obtain your **install-config.yaml** file by using the installation program and modify it.

```

apiVersion: v1
baseDomain: example.com
compute: ① ②
- architecture: ppc64le
  hyperthreading: Enabled ③
  name: worker
  platform:
    powervs:
      smtLevel: 8 ④
  replicas: 3
controlPlane: ⑤ ⑥
  architecture: ppc64le
  hyperthreading: Enabled ⑦
  name: master
  platform:
    powervs:
      smtLevel: 8 ⑧
  replicas: 3
metadata:
  creationTimestamp: null
  name: example-cluster-existing-vpc

```

```

networking:
  clusterNetwork:
    - cidr: 10.128.0.0/14 9
      hostPrefix: 23
  machineNetwork:
    - cidr: 192.168.0.0/24
  networkType: OVNKubernetes 10
  serviceNetwork:
    - 172.30.0.0/16
platform:
  powervs:
    userID: ibm-user-id
    powervsResourceGroup: "ibmcloud-resource-group"
    region: powervs-region
    vpcRegion : vpc-region
    vpcName: name-of-existing-vpc 11
    vpcSubnets: 12
    - powervs-region-example-subnet-1
    zone: powervs-zone
    serviceInstanceGUID: "powervs-region-service-instance-guid"
  credentialsMode: Manual
  publish: External 13
  pullSecret: '{"auths": ...}' 14
  fips: false
  sshKey: ssh-ed25519 AAAA... 15

```

- 1 5 If you do not provide these parameters and values, the installation program provides the default value.
- 2 6 The **controlPlane** section is a single mapping, but the compute section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Both sections currently define a single machine pool. Only one control plane pool is used.
- 3 7 Whether to enable or disable simultaneous multithreading, or **hyperthreading**. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores. You can disable it by setting the parameter value to **Disabled**. If you disable simultaneous multithreading in some cluster machines, you must disable it in all cluster machines.
- 4 8 The `smtLevel` specifies the level of SMT to set to the control plane and compute machines. The supported values are 1, 2, 4, 8, **'off'** and **'on'**. The default value is 8. The `smtLevel` **'off'** sets SMT to off and `smtLevel` **'on'** sets SMT to the default value 8 on the cluster nodes.



NOTE

When simultaneous multithreading (SMT), or hyperthreading is not enabled, one vCPU is equivalent to one physical core. When enabled, total vCPUs is computed as (Thread(s) per core * Core(s) per socket) * Socket(s). The `smtLevel` controls the threads per core. Lower SMT levels may require additional assigned cores when deploying the cluster nodes. You can do this by setting the **'processors'** parameter in the **install-config.yaml** file to an appropriate value to meet the requirements for deploying OpenShift Container Platform successfully.

- 9 The machine CIDR must contain the subnets for the compute machines and control plane machines.

- 10 The cluster network plugin for installation. The supported value is **OVNKubernetes**.
- 11 Specify the name of an existing VPC.
- 12 Specify the name of the existing VPC subnet. The subnets must belong to the VPC that you specified. Specify a subnet for each availability zone in the region.
- 13 Specify how to publish the user-facing endpoints of your cluster.
- 14 Required. The installation program prompts you for this value.
- 15 Provide the **sshKey** value that you use to access the machines in your cluster.



IMPORTANT

If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

20.5.7.3. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```

apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5

```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.



NOTE

The installation program does not support the proxy **readinessEndpoints** field.



NOTE

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.



NOTE

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

20.5.8. Manually creating IAM

Installing the cluster requires that the Cloud Credential Operator (CCO) operate in manual mode. While the installation program configures the CCO for manual mode, you must specify the identity and access management secrets for your cloud provider.

You can use the Cloud Credential Operator (CCO) utility (**ccoctl**) to create the required IBM Cloud® resources.

Prerequisites

- You have configured the **ccoctl** binary.
- You have an existing **install-config.yaml** file.

Procedure

1. Edit the **install-config.yaml** configuration file so that it contains the **credentialsMode** parameter set to **Manual**.

Example **install-config.yaml** configuration file

```
apiVersion: v1
baseDomain: cluster1.example.com
credentialsMode: Manual 1
compute:
- architecture: ppc64le
  hyperthreading: Enabled
```

- 1** This line is added to set the **credentialsMode** parameter to **Manual**.

2. To generate the manifests, run the following command from the directory that contains the installation program:

```
$ openshift-install create manifests --dir <installation_directory>
```

3. From the directory that contains the installation program, set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

4. Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included \ 1
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml \ 2
  --to=<path_to_directory_for_credentials_requests> \ 3
```

- 1 The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2 Specify the location of the **install-config.yaml** file.
- 3 Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

This command creates a YAML file for each **CredentialsRequest** object.

Sample CredentialsRequest object

```
apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  labels:
    controller-tools.k8s.io: "1.0"
  name: openshift-image-registry-ibmcos
  namespace: openshift-cloud-credential-operator
spec:
  secretRef:
    name: installer-cloud-credentials
    namespace: openshift-image-registry
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: IBMCloudProviderSpec
    policies:
      - attributes:
          - name: serviceName
            value: cloud-object-storage
        roles:
          - crn:v1:bluemix:public:iam::::role:Viewer
          - crn:v1:bluemix:public:iam::::role:Operator
          - crn:v1:bluemix:public:iam::::role:Editor
          - crn:v1:bluemix:public:iam::::serviceRole:Reader
          - crn:v1:bluemix:public:iam::::serviceRole:Writer
      - attributes:
          - name: resourceType
            value: resource-group
        roles:
          - crn:v1:bluemix:public:iam::::role:Viewer
```

5. Create the service ID for each credential request, assign the policies defined, create an API key, and generate the secret:

```
$ ccoctl ibmcloud create-service-id \
  --credentials-requests-dir=<path_to_credential_requests_directory> \ 1
```

```
--name=<cluster_name> \ 2
--output-dir=<installation_directory> \ 3
--resource-group-name=<resource_group_name> 4
```

- 1 Specify the directory containing the files for the component **CredentialsRequest** objects.
- 2 Specify the name of the OpenShift Container Platform cluster.
- 3 Optional: Specify the directory in which you want the **ccoctl** utility to create objects. By default, the utility creates objects in the directory in which the commands are run.
- 4 Optional: Specify the name of the resource group used for scoping the access policies.



NOTE

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

If an incorrect resource group name is provided, the installation fails during the bootstrap phase. To find the correct resource group name, run the following command:

```
$ grep resourceGroup <installation_directory>/manifests/cluster-infrastructure-02-config.yml
```

Verification

- Ensure that the appropriate secrets were generated in your cluster's **manifests** directory.

20.5.9. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.

Procedure

- Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

- 1 For **<installation_directory>**, specify the location of your customized **./install-config.yaml** file.
- 2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to **<installation_directory>/openshift_install.log**.



IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

20.5.10. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.



IMPORTANT

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**.
To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

■

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.



NOTE

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your PATH.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

20.5.11. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

Additional resources

- [Accessing the web console](#)

20.5.12. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- [About remote health monitoring](#)

20.5.13. Next steps

- [Customize your cluster](#)
- Optional: [Opt out of remote health reporting](#)

20.6. INSTALLING A PRIVATE CLUSTER ON IBM POWER VIRTUAL SERVER

In OpenShift Container Platform version 4.16, you can install a private cluster into an existing VPC and IBM Power® Virtual Server Workspace. The installation program provisions the rest of the required infrastructure, which you can further customize. To customize the installation, you modify parameters in the **install-config.yaml** file before you install the cluster.

20.6.1. Prerequisites


- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.

- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [configured an IBM Cloud® account](#) to host the cluster.
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.
- You configured the **ccoctl** utility before you installed the cluster. For more information, see [Configuring the Cloud Credential Operator utility](#).

20.6.2. Private clusters

You can deploy a private OpenShift Container Platform cluster that does not expose external endpoints. Private clusters are accessible from only an internal network and are not visible to the internet.

By default, OpenShift Container Platform is provisioned to use publicly-accessible DNS and endpoints. A private cluster sets the DNS, Ingress Controller, and API server to private when you deploy your cluster. This means that the cluster resources are only accessible from your internal network and are not visible to the internet.



IMPORTANT

If the cluster has any public subnets, load balancer services created by administrators might be publicly accessible. To ensure cluster security, verify that these services are explicitly annotated as private.

To deploy a private cluster, you must:

- Use existing networking that meets your requirements.
- Create a DNS zone using IBM Cloud® DNS Services and specify it as the base domain of the cluster. For more information, see "Using IBM Cloud® DNS Services to configure DNS resolution".
- Deploy from a machine that has access to:
 - The API services for the cloud to which you provision.
 - The hosts on the network that you provision.
 - The internet to obtain installation media.

You can use any machine that meets these access requirements and follows your company's guidelines. For example, this machine can be a bastion host on your cloud network or a machine that has access to the network through a VPN.

20.6.3. Private clusters in IBM Power Virtual Server

To create a private cluster on IBM Power® Virtual Server, you must provide an existing private Virtual Private Cloud (VPC) and subnets to host the cluster. The installation program must also be able to resolve the DNS records that the cluster requires. The installation program configures the Ingress Operator and API server for only internal traffic.

The cluster still requires access to internet to access the IBM Cloud® APIs.

The following items are not required or created when you install a private cluster:

- Public subnets
- Public network load balancers, which support public Ingress
- A public DNS zone that matches the **baseDomain** for the cluster

You will also need to create an IBM® DNS service containing a DNS zone that matches your **baseDomain**. Unlike standard deployments on Power VS which use IBM® CIS for DNS, you must use IBM® DNS for your DNS service.

20.6.3.1. Limitations

Private clusters on IBM Power® Virtual Server are subject only to the limitations associated with the existing VPC that was used for cluster deployment.

20.6.4. Requirements for using your VPC

You must correctly configure the existing VPC and its subnets before you install the cluster. The installation program does not create a VPC or VPC subnet in this scenario.

The installation program cannot:

- Subdivide network ranges for the cluster to use
- Set route tables for the subnets
- Set VPC options like DHCP



NOTE

The installation program requires that you use the cloud-provided DNS server. Using a custom DNS server is not supported and causes the installation to fail.

20.6.4.1. VPC validation

The VPC and all of the subnets must be in an existing resource group. The cluster is deployed to this resource group.

As part of the installation, specify the following in the **install-config.yaml** file:

- The name of the resource group
- The name of VPC
- The name of the VPC subnet

To ensure that the subnets that you provide are suitable, the installation program confirms that all of the subnets you specify exists.



NOTE

Subnet IDs are not supported.

20.6.4.2. Isolation between clusters

If you deploy OpenShift Container Platform to an existing network, the isolation of cluster services is reduced in the following ways:


- ICMP Ingress is allowed to the entire network.
- TCP port 22 Ingress (SSH) is allowed to the entire network.
- Control plane TCP 6443 Ingress (Kubernetes API) is allowed to the entire network.
- Control plane TCP 22623 Ingress (MCS) is allowed to the entire network.

20.6.5. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.



IMPORTANT

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

20.6.6. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.

**NOTE**

You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.

**NOTE**

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the `ssh-agent` process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```

4. Add your SSH private key to the `ssh-agent`:

```
$ ssh-add <path>/<file_name> 1
```

- 1 Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

20.6.7. Obtaining the installation program


Before you install OpenShift Container Platform, download the installation file on the host you are using for installation.

Prerequisites

- You have a computer that runs Linux or macOS, with 500 MB of local disk space.


Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select your infrastructure provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.



IMPORTANT

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.



IMPORTANT

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.

4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#). This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

20.6.8. Exporting the API key

You must set the API key you created as a global variable; the installation program ingests the variable during startup to set the API key.

Prerequisites

- You have created either a user API key or service ID API key for your IBM Cloud® account.

Procedure

- Export your API key for your account as a global variable:

```
$ export IBMCLLOUD_API_KEY=<api_key>
```



IMPORTANT

You must set the variable name exactly as specified; the installation program expects the variable name to be present during startup.

20.6.9. Manually creating the installation configuration file

Installing the cluster requires that you manually create the installation configuration file.

Prerequisites

- You have an SSH public key on your local machine to provide to the installation program. The key will be used for SSH authentication onto your cluster nodes for debugging and disaster recovery.
- You have obtained the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create an installation directory to store your required installation assets in:

```
$ mkdir <installation_directory>
```



IMPORTANT

You must create a directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.


2. Customize the sample **install-config.yaml** file template that is provided and save it in the **<installation_directory>**.



NOTE

You must name this configuration file **install-config.yaml**.

3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the next step of the installation process. You must back it up now.

Additional resources

- [Installation configuration parameters for IBM Power® Virtual Server](#)


20.6.9.1. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 20.4. Minimum resource requirements

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Bootstrap	RHCOS	2	16 GB	100 GB	300
Control plane	RHCOS	2	16 GB	100 GB	300
Compute	RHCOS	2	8 GB	100 GB	300

1. One vCPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio: $(\text{threads per core} \times \text{cores}) \times \text{sockets} = \text{vCPUs}$.
2. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes. Note that on many cloud platforms, storage size and IOPS scale together, so you might need to over-allocate storage volume to obtain sufficient performance.



NOTE

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

20.6.9.2. Sample customized install-config.yaml file for IBM Power Virtual Server

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.



IMPORTANT

This sample YAML file is provided for reference only. You must obtain your **install-config.yaml** file by using the installation program and modify it.

```

apiVersion: v1
baseDomain: example.com
compute: ① ②
- architecture: ppc64le
  hyperthreading: Enabled ③
  name: worker
  platform:
    powervs:
      smtLevel: 8 ④
  replicas: 3
controlPlane: ⑤ ⑥
  architecture: ppc64le
  hyperthreading: Enabled ⑦
  name: master
  platform:
    powervs:
      smtLevel: 8 ⑧
  replicas: 3
metadata:
  creationTimestamp: null
  name: example-private-cluster-name
networking:
  clusterNetwork:
  - cidr: 10.128.0.0/14 ⑨
    hostPrefix: 23
  machineNetwork:
  - cidr: 192.168.0.0/24
  networkType: OVNKubernetes ⑩
  serviceNetwork:
  - 172.30.0.0/16
platform:
  powervs:
    userID: ibm-user-id
    powervsResourceGroup: "ibmcloud-resource-group"
    region: powervs-region
    vpcName: name-of-existing-vpc ⑪

```

```

vpcSubnets:
- powervs-region-example-subnet-1
vpcRegion : vpc-region
zone: powervs-zone
serviceInstanceGUID: "powervs-region-service-instance-guid"
publish: Internal 12
pullSecret: '{"auths": ...}' 13
sshKey: ssh-ed25519 AAAA... 14

```

- 1 5** If you do not provide these parameters and values, the installation program provides the default value.
- 2 6** The **controlPlane** section is a single mapping, but the compute section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Both sections currently define a single machine pool. Only one control plane pool is used.
- 3 7** Whether to enable or disable simultaneous multithreading, or **hyperthreading**. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores. You can disable it by setting the parameter value to **Disabled**. If you disable simultaneous multithreading in some cluster machines, you must disable it in all cluster machines.
- 4 8** The `smtLevel` specifies the level of SMT to set to the control plane and compute machines. The supported values are 1, 2, 4, 8, **'off'** and **'on'**. The default value is 8. The `smtLevel` **'off'** sets SMT to off and `smtlevel` **'on'** sets SMT to the default value 8 on the cluster nodes.



NOTE

When simultaneous multithreading (SMT), or hyperthreading is not enabled, one vCPU is equivalent to one physical core. When enabled, total vCPUs is computed as (Thread(s) per core * Core(s) per socket) * Socket(s). The `smtLevel` controls the threads per core. Lower SMT levels may require additional assigned cores when deploying the cluster nodes. You can do this by setting the **'processors'** parameter in the **install-config.yaml** file to an appropriate value to meet the requirements for deploying OpenShift Container Platform successfully.

- 9** The machine CIDR must contain the subnets for the compute machines and control plane machines.
- 10** The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 11** Specify the name of an existing VPC.
- 12** Specify how to publish the user-facing endpoints of your cluster. Set `publish` to **Internal** to deploy a private cluster.
- 13** Required. The installation program prompts you for this value.
- 14** Provide the **sshKey** value that you use to access the machines in your cluster.



IMPORTANT

If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.

**NOTE**

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

20.6.9.3. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.

**NOTE**

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
additionalTrustBundle: | 4
  -----BEGIN CERTIFICATE-----
  <MY_TRUSTED_CA_CERT>
  -----END CERTIFICATE-----
additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.

- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with `.` to match subdomains only. For
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.

**NOTE**

The installation program does not support the proxy **readinessEndpoints** field.

**NOTE**

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.

**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

20.6.10. Manually creating IAM

Installing the cluster requires that the Cloud Credential Operator (CCO) operate in manual mode. While the installation program configures the CCO for manual mode, you must specify the identity and access management secrets for your cloud provider.

You can use the Cloud Credential Operator (CCO) utility (**ccoctl**) to create the required IBM Cloud® resources.

Prerequisites

- You have configured the **ccoctl** binary.
- You have an existing **install-config.yaml** file.

Procedure

1. Edit the **install-config.yaml** configuration file so that it contains the **credentialsMode** parameter set to **Manual**.

Example install-config.yaml configuration file

```
apiVersion: v1
baseDomain: cluster1.example.com
credentialsMode: Manual 1
compute:
- architecture: ppc64le
  hyperthreading: Enabled
```

- 1** This line is added to set the **credentialsMode** parameter to **Manual**.

2. To generate the manifests, run the following command from the directory that contains the installation program:

```
$ openshift-install create manifests --dir <installation_directory>
```

3. From the directory that contains the installation program, set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

4. Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included 1 \
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml 2 \
  --to=<path_to_directory_for_credentials_requests> 3
```

- 1** The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2** Specify the location of the **install-config.yaml** file.
- 3** Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

This command creates a YAML file for each **CredentialsRequest** object.

Sample CredentialsRequest object

```
apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  labels:
```

```

    controller-tools.k8s.io: "1.0"
    name: openshift-image-registry-ibmcos
    namespace: openshift-cloud-credential-operator
spec:
  secretRef:
    name: installer-cloud-credentials
    namespace: openshift-image-registry
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: IBMCloudProviderSpec
    policies:
      - attributes:
          - name: serviceName
            value: cloud-object-storage
        roles:
          - crn:v1:bluemix:public:iam::::role:Viewer
          - crn:v1:bluemix:public:iam::::role:Operator
          - crn:v1:bluemix:public:iam::::role:Editor
          - crn:v1:bluemix:public:iam::::serviceRole:Reader
          - crn:v1:bluemix:public:iam::::serviceRole:Writer
      - attributes:
          - name: resourceType
            value: resource-group
        roles:
          - crn:v1:bluemix:public:iam::::role:Viewer

```

5. Create the service ID for each credential request, assign the policies defined, create an API key, and generate the secret:

```

$ ccoctl ibmcloud create-service-id \
  --credentials-requests-dir=<path_to_credential_requests_directory> \1
  --name=<cluster_name> \2
  --output-dir=<installation_directory> \3
  --resource-group-name=<resource_group_name> \4

```

- 1 Specify the directory containing the files for the component **CredentialsRequest** objects.
- 2 Specify the name of the OpenShift Container Platform cluster.
- 3 Optional: Specify the directory in which you want the **ccoctl** utility to create objects. By default, the utility creates objects in the directory in which the commands are run.
- 4 Optional: Specify the name of the resource group used for scoping the access policies.



NOTE

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

If an incorrect resource group name is provided, the installation fails during the bootstrap phase. To find the correct resource group name, run the following command:


```
$ grep resourceGroup <installation_directory>/manifests/cluster-infrastructure-02-config.yml
```

Verification

- Ensure that the appropriate secrets were generated in your cluster's **manifests** directory.

20.6.11. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.

Procedure

- Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1  
--log-level=info 2
```

- 1** For **<installation_directory>**, specify the location of your customized **./install-config.yaml** file.
- 2** To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to **<installation_directory>/openshift_install.log**.



IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

20.6.12. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.



IMPORTANT

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**.
To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:


```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.

**NOTE**

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your PATH.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

20.6.13. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
-
```

```
system:admin
```

Additional resources

- [Accessing the web console](#)

20.6.14. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- [About remote health monitoring](#)

20.6.15. Next steps

- [Customize your cluster](#)
- Optional: [Opt out of remote health reporting](#)

20.7. INSTALLING A CLUSTER ON IBM POWER VIRTUAL SERVER IN A RESTRICTED NETWORK

In OpenShift Container Platform 4.16, you can install a cluster on IBM Cloud® in a restricted network by creating an internal mirror of the installation release content on an existing Virtual Private Cloud (VPC) on IBM Cloud®.

20.7.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [configured an IBM Cloud® account](#) to host the cluster.
- You [mirrored the images for a disconnected installation](#) to your registry and obtained the **imageContentSources** data for your version of OpenShift Container Platform.



IMPORTANT

Because the installation media is on the mirror host, you can use that computer to complete all installation steps.

- You have an existing VPC in IBM Cloud®. When installing a cluster in a restricted network, you cannot use the installer-provisioned VPC. You must use a user-provisioned VPC that satisfies one of the following requirements:
 - Contains the mirror registry
 - Has firewall rules or a peering connection to access the mirror registry hosted elsewhere
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.
- You configured the **ccoctl** utility before you installed the cluster. For more information, see [Configuring the Cloud Credential Operator utility](#).

20.7.2. About installations in restricted networks

In OpenShift Container Platform 4.16, you can perform an installation that does not require an active connection to the internet to obtain software components. Restricted network installations can be completed using installer-provisioned infrastructure or user-provisioned infrastructure, depending on the cloud platform to which you are installing the cluster.

If you choose to perform a restricted network installation on a cloud platform, you still require access to its cloud APIs. Some cloud functions, like Amazon Web Service's Route 53 DNS and IAM services, require internet access. Depending on your network, you might require less internet access for an installation on bare metal hardware, Nutanix, or on VMware vSphere.

To complete a restricted network installation, you must create a registry that mirrors the contents of the OpenShift image registry and contains the installation media. You can create this registry on a mirror host, which can access both the internet and your closed network, or by using other methods that meet your restrictions.

20.7.2.1. Additional limits

Clusters in restricted networks have the following additional limitations and restrictions:

- The **ClusterVersion** status includes an **Unable to retrieve available updates** error.
- By default, you cannot use the contents of the Developer Catalog because you cannot access the required image stream tags.

20.7.3. About using a custom VPC

In OpenShift Container Platform 4.16, you can deploy a cluster into the subnets of an existing IBM® Virtual Private Cloud (VPC).

20.7.3.1. Requirements for using your VPC

You must correctly configure the existing VPC and its subnets before you install the cluster. The installation program does not create a VPC or VPC subnet in this scenario.

The installation program cannot:

- Subdivide network ranges for the cluster to use
- Set route tables for the subnets
- Set VPC options like DHCP

**NOTE**

The installation program requires that you use the cloud-provided DNS server. Using a custom DNS server is not supported and causes the installation to fail.


20.7.3.2. VPC validation

The VPC and all of the subnets must be in an existing resource group. The cluster is deployed to this resource group.

As part of the installation, specify the following in the **install-config.yaml** file:

- The name of the resource group
- The name of VPC
- The name of the VPC subnet

To ensure that the subnets that you provide are suitable, the installation program confirms that all of the subnets you specify exists.

**NOTE**

Subnet IDs are not supported.

20.7.3.3. Isolation between clusters

If you deploy OpenShift Container Platform to an existing network, the isolation of cluster services is reduced in the following ways:

- ICMP Ingress is allowed to the entire network.
- TCP port 22 Ingress (SSH) is allowed to the entire network.
- Control plane TCP 6443 Ingress (Kubernetes API) is allowed to the entire network.
- Control plane TCP 22623 Ingress (MCS) is allowed to the entire network.

20.7.4. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to obtain the images that are necessary to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.

20.7.5. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.



NOTE

You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.



NOTE

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

20.7.6. Exporting the API key

You must set the API key you created as a global variable; the installation program ingests the variable during startup to set the API key.

Prerequisites

- You have created either a user API key or service ID API key for your IBM Cloud® account.

Procedure

- Export your API key for your account as a global variable:

```
$ export IBM_CLOUD_API_KEY=<api_key>
```



IMPORTANT

You must set the variable name exactly as specified; the installation program expects the variable name to be present during startup.

20.7.7. Creating the installation configuration file

You can customize the OpenShift Container Platform cluster you install on

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster. For a restricted network installation, these files are on your mirror host.
- You have the **imageContentSources** values that were generated during mirror registry creation.
- You have obtained the contents of the certificate for your mirror registry.
- You have retrieved a Red Hat Enterprise Linux CoreOS (RHCOS) image and uploaded it to an accessible location.

Procedure

1. Create the **install-config.yaml** file.
 - a. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install create install-config --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the directory name to store the files that the installation program creates.

When specifying the directory:

- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
- Use an empty directory. Some installation assets, such as bootstrap X.509 certificates, have short expiration intervals, therefore you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.



NOTE

Always delete the **~/powervs** directory to avoid reusing a stale configuration. Run the following command:

```
$ rm -rf ~/.powervs
```

- b. At the prompts, provide the configuration details for your cloud:
 - i. Optional: Select an SSH key to use to access your cluster machines.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- ii. Select **powervs** as the platform to target.

- iii. Select the region to deploy the cluster to.
 - iv. Select the zone to deploy the cluster to.
 - v. Select the base domain to deploy the cluster to. The base domain corresponds to the public DNS zone that you created for your cluster.
 - vi. Enter a descriptive name for your cluster.
2. Edit the **install-config.yaml** file to give the additional information that is required for an installation in a restricted network.

- a. Update the **pullSecret** value to contain the authentication information for your registry:

```
pullSecret: '{"auths":{"<mirror_host_name>:5000": {"auth": "<credentials>","email": "you@example.com"}}}'
```

For **<mirror_host_name>**, specify the registry domain name that you specified in the certificate for your mirror registry, and for **<credentials>**, specify the base64-encoded user name and password for your mirror registry.

- b. Add the **additionalTrustBundle** parameter and value.

```
additionalTrustBundle: |
  -----BEGIN CERTIFICATE-----
  ////////////////////////////////////////////////////////////////////
  -----END CERTIFICATE-----
```

The value must be the contents of the certificate file that you used for your mirror registry. The certificate file can be an existing, trusted certificate authority, or the self-signed certificate that you generated for the mirror registry.

- c. Define the network and subnets for the VPC to install the cluster in under the parent **platform.ibmcloud** field:

```
vpcName: <existing_vpc>
vpcSubnets: <vpcSubnet>
```

For **platform.powervs.vpcName**, specify the name for the existing IBM Cloud®. For **platform.powervs.vpcSubnets**, specify the existing subnets.

- d. Add the image content resources, which resemble the following YAML excerpt:

```
imageContentSources:
- mirrors:
  - <mirror_host_name>:5000/<repo_name>/release
  source: quay.io/openshift-release-dev/ocp-release
- mirrors:
  - <mirror_host_name>:5000/<repo_name>/release
  source: registry.redhat.io/ocp/release
```


For these values, use the **imageContentSources** that you recorded during mirror registry creation.

- e. Optional: Set the publishing strategy to **Internal**:

publish: Internal

By setting this option, you create an internal Ingress Controller and a private load balancer.

3. Make any other modifications to the **install-config.yaml** file that you require.
For more information about the parameters, see "Installation configuration parameters".
4. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

Additional resources

- [Installation configuration parameters for IBM Power® Virtual Server](#)

20.7.7.1. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 20.5. Minimum resource requirements

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Bootstrap	RHCOS	2	16 GB	100 GB	300
Control plane	RHCOS	2	16 GB	100 GB	300
Compute	RHCOS	2	8 GB	100 GB	300

1. One vCPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio: (threads per core × cores) × sockets = vCPUs.
2. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes. Note that on many cloud platforms, storage size and IOPS scale together, so you might need to over-allocate storage volume to obtain sufficient performance.

**NOTE**

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

20.7.7.2. Sample customized install-config.yaml file for IBM Power Virtual Server

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.

**IMPORTANT**

This sample YAML file is provided for reference only. You must obtain your **install-config.yaml** file by using the installation program and modify it.

```

apiVersion: v1
baseDomain: example.com 1
controlPlane: 2 3
  hyperthreading: Enabled 4
  name: master
  platform:
    powervs:
      smtLevel: 8 5
  replicas: 3
compute: 6 7
- hyperthreading: Enabled 8
  name: worker
  platform:
    powervs:
      smtLevel: 8 9
    ibmcloud: {}
  replicas: 3
metadata:
  name: example-restricted-cluster-name 10
networking:

```

```

clusterNetwork:
- cidr: 10.128.0.0/14 11
  hostPrefix: 23
machineNetwork:
- cidr: 10.0.0.0/16 12
networkType: OVNKubernetes 13
serviceNetwork:
- 192.168.0.0/24
platform:
powervs:
  userid: ibm-user-id
  powervsResourceGroup: "ibmcloud-resource-group" 14
  region: "powervs-region"
  vpcRegion: "vpc-region"
  vpcName: name-of-existing-vpc 15
  vpcSubnets: 16
    - name-of-existing-vpc-subnet
  zone: "powervs-zone"
  serviceInstanceID: "service-instance-id"
publish: Internal
credentialsMode: Manual
pullSecret: '{"auths":{"<local_registry>":{"auth": "<credentials>","email": "you@example.com"}}}' 17
sshKey: ssh-ed25519 AAAA... 18
additionalTrustBundle: | 19
  -----BEGIN CERTIFICATE-----
  <MY_TRUSTED_CA_CERT>
  -----END CERTIFICATE-----
imageContentSources: 20
- mirrors:
  - <local_registry>/<local_repository_name>/release
  source: quay.io/openshift-release-dev/ocp-release
- mirrors:
  - <local_registry>/<local_repository_name>/release
  source: quay.io/openshift-release-dev/ocp-v4.0-art-dev

```

1 10 Required.

2 6 If you do not provide these parameters and values, the installation program provides the default value.

3 7 The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.

4 8 Enables or disables simultaneous multithreading, also known as Hyper-Threading. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores. You can disable it by setting the parameter value to **Disabled**. If you disable simultaneous multithreading in some cluster machines, you must disable it in all cluster machines.



IMPORTANT

If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance. Use larger machine types, such as **n1-standard-8**, for your machines if you disable simultaneous multithreading.

- 5 9** The `smtLevel` specifies the level of SMT to set to the control plane and compute machines. The supported values are 1, 2, 4, 8, **'off'** and **'on'**. The default value is 8. The `smtLevel` **'off'** sets SMT to off and `smtlevel` **'on'** sets SMT to the default value 8 on the cluster nodes.



NOTE

When simultaneous multithreading (SMT), or hyperthreading is not enabled, one vCPU is equivalent to one physical core. When enabled, total vCPUs is computed as (Thread(s) per core * Core(s) per socket) * Socket(s). The `smtLevel` controls the threads per core. Lower SMT levels may require additional assigned cores when deploying the cluster nodes. You can do this by setting the **'processors'** parameter in the **install-config.yaml** file to an appropriate value to meet the requirements for deploying OpenShift Container Platform successfully.

- 11** The machine CIDR must contain the subnets for the compute machines and control plane machines.
- 12** The CIDR must contain the subnets defined in **platform.ibmcloud.controlPlaneSubnets** and **platform.ibmcloud.computeSubnets**.
- 13** The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 14** The name of an existing resource group. The existing VPC and subnets should be in this resource group. The cluster is deployed to this resource group.
- 15** Specify the name of an existing VPC.
- 16** Specify the name of the existing VPC subnet. The subnets must belong to the VPC that you specified. Specify a subnet for each availability zone in the region.
- 17** For **<local_registry>**, specify the registry domain name, and optionally the port, that your mirror registry uses to serve content. For example, `registry.example.com` or `registry.example.com:5000`. For **<credentials>**, specify the base64-encoded user name and password for your mirror registry.
- 18** You can optionally provide the **sshKey** value that you use to access the machines in your cluster.
- 19** Provide the contents of the certificate file that you used for your mirror registry.
- 20** Provide the **imageContentSources** section from the output of the command to mirror the repository.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

20.7.7.3. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then

creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.

- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.



NOTE

The installation program does not support the proxy **readinessEndpoints** field.



NOTE

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.



NOTE

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

20.7.8. Manually creating IAM

Installing the cluster requires that the Cloud Credential Operator (CCO) operate in manual mode. While the installation program configures the CCO for manual mode, you must specify the identity and access management secrets for your cloud provider.

You can use the Cloud Credential Operator (CCO) utility (**ccoctl**) to create the required IBM Cloud® resources.

Prerequisites

- You have configured the **ccoctl** binary.
- You have an existing **install-config.yaml** file.

Procedure

1. Edit the **install-config.yaml** configuration file so that it contains the **credentialsMode** parameter set to **Manual**.

Example `install-config.yaml` configuration file

```
apiVersion: v1
baseDomain: cluster1.example.com
credentialsMode: Manual 1
compute:
- architecture: ppc64le
  hyperthreading: Enabled
```

- 1** This line is added to set the **credentialsMode** parameter to **Manual**.
- To generate the manifests, run the following command from the directory that contains the installation program:

```
$ openshift-install create manifests --dir <installation_directory>
```

- From the directory that contains the installation program, set a **\$RELEASE_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

- Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --included 1 \
  --install-config=<path_to_directory_with_installation_configuration>/install-config.yaml 2 \
  --to=<path_to_directory_for_credentials_requests> 3
```

- 1** The **--included** parameter includes only the manifests that your specific cluster configuration requires.
- 2** Specify the location of the **install-config.yaml** file.
- 3** Specify the path to the directory where you want to store the **CredentialsRequest** objects. If the specified directory does not exist, this command creates it.

This command creates a YAML file for each **CredentialsRequest** object.

Sample **CredentialsRequest** object

```
apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  labels:
    controller-tools.k8s.io: "1.0"
  name: openshift-image-registry-ibmcos
  namespace: openshift-cloud-credential-operator
spec:
  secretRef:
```

```

name: installer-cloud-credentials
namespace: openshift-image-registry
providerSpec:
  apiVersion: cloudcredential.openshift.io/v1
  kind: IBMCloudProviderSpec
  policies:
  - attributes:
    - name: serviceName
      value: cloud-object-storage
    roles:
    - crn:v1:bluemix:public:iam::::role:Viewer
    - crn:v1:bluemix:public:iam::::role:Operator
    - crn:v1:bluemix:public:iam::::role:Editor
    - crn:v1:bluemix:public:iam::::serviceRole:Reader
    - crn:v1:bluemix:public:iam::::serviceRole:Writer
  - attributes:
    - name: resourceType
      value: resource-group
    roles:
    - crn:v1:bluemix:public:iam::::role:Viewer

```

5. Create the service ID for each credential request, assign the policies defined, create an API key, and generate the secret:

```

$ ccoctl ibmcloud create-service-id \
  --credentials-requests-dir=<path_to_credential_requests_directory> \ 1
  --name=<cluster_name> \ 2
  --output-dir=<installation_directory> \ 3
  --resource-group-name=<resource_group_name> \ 4

```

- 1** Specify the directory containing the files for the component **CredentialsRequest** objects.
- 2** Specify the name of the OpenShift Container Platform cluster.
- 3** Optional: Specify the directory in which you want the **ccoctl** utility to create objects. By default, the utility creates objects in the directory in which the commands are run.
- 4** Optional: Specify the name of the resource group used for scoping the access policies.



NOTE

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

If an incorrect resource group name is provided, the installation fails during the bootstrap phase. To find the correct resource group name, run the following command:

```

$ grep resourceGroup <installation_directory>/manifests/cluster-infrastructure-02-config.yml


```

Verification

- Ensure that the appropriate secrets were generated in your cluster's **manifests** directory.

20.7.9. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have configured an account with the cloud platform that hosts your cluster.
- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.

Procedure

- Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

- 1 For **<installation_directory>**, specify the location of your customized **./install-config.yaml** file.
- 2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to **<installation_directory>/openshift_install.log**.



IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
```

```
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

20.7.10. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.



IMPORTANT

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**.
To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.



NOTE

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your PATH.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

20.7.11. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

Additional resources

- [Accessing the web console](#)

20.7.12. Disabling the default OperatorHub catalog sources

Operator catalogs that source content provided by Red Hat and community projects are configured for OperatorHub by default during an OpenShift Container Platform installation. In a restricted network environment, you must disable the default catalogs as a cluster administrator.

Procedure

- Disable the sources for the default catalogs by adding **disableAllDefaultSources: true** to the **OperatorHub** object:

```
$ oc patch OperatorHub cluster --type json \
-p '[{"op": "add", "path": "/spec/disableAllDefaultSources", "value": true}]'
```

TIP

Alternatively, you can use the web console to manage catalog sources. From the **Administration** → **Cluster Settings** → **Configuration** → **OperatorHub** page, click the **Sources** tab, where you can create, update, delete, disable, and enable individual sources.

20.7.13. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- [About remote health monitoring](#)

20.7.14. Next steps

- [Customize your cluster](#)
- Optional: [Opt out of remote health reporting](#)
- Optional: [Registering your disconnected cluster](#)

20.8. UNINSTALLING A CLUSTER ON IBM POWER VIRTUAL SERVER

You can remove a cluster that you deployed to IBM Power® Virtual Server.

20.8.1. Removing a cluster that uses installer-provisioned infrastructure

You can remove a cluster that uses installer-provisioned infrastructure from your cloud.



NOTE

After uninstallation, check your cloud provider for any resources not removed properly, especially with User Provisioned Infrastructure (UPI) clusters. There might be resources that the installer did not create or that the installer is unable to access.

Prerequisites

- You have a copy of the installation program that you used to deploy the cluster.
- You have the files that the installation program generated when you created your cluster.

- You have configured the **ccoctl** binary.
- You have installed the IBM Cloud® CLI and installed or updated the VPC infrastructure service plugin. For more information see "Prerequisites" in the [IBM Cloud® CLI documentation](#).

Procedure

1. If the following conditions are met, this step is required:

- The installer created a resource group as part of the installation process.
- You or one of your applications created persistent volume claims (PVCs) after the cluster was deployed.

In which case, the PVCs are not removed when uninstalling the cluster, which might prevent the resource group from being successfully removed. To prevent a failure:

- Log in to the IBM Cloud® using the CLI.
- To list the PVCs, run the following command:

```
$ ibmcloud is volumes --resource-group-name <infrastructure_id>
```

For more information about listing volumes, see the [IBM Cloud® CLI documentation](#).

- To delete the PVCs, run the following command:

```
$ ibmcloud is volume-delete --force <volume_id>
```

For more information about deleting volumes, see the [IBM Cloud® CLI documentation](#).

2. Export the API key that was created as part of the installation process.

```
$ export IBM_CLOUD_API_KEY=<api_key>
```



NOTE

You must set the variable name exactly as specified. The installation program expects the variable name to be present to remove the service IDs that were created when the cluster was installed.

3. From the directory that contains the installation program on the computer that you used to install the cluster, run the following command:

```
$ ./openshift-install destroy cluster \
--dir <installation_directory> --log-level info 1 2
```

1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2 To view different details, specify **warn**, **debug**, or **error** instead of **info**.

**NOTE**

- You must specify the directory that contains the cluster definition files for your cluster. The installation program requires the **metadata.json** file in this directory to delete the cluster.
- You might have to run the **openshift-install destroy** command up to three times to ensure a proper cleanup.

4. Remove the manual CCO credentials that were created for the cluster:

```
$ ccoctl ibmcloud delete-service-id \
  --credentials-requests-dir <path_to_credential_requests_directory> \
  --name <cluster_name>
```

**NOTE**

If your cluster uses Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set, you must include the **--enable-tech-preview** parameter.

5. Optional: Delete the **<installation_directory>** directory and the OpenShift Container Platform installation program.

20.9. INSTALLATION CONFIGURATION PARAMETERS FOR IBM POWER VIRTUAL SERVER

Before you deploy an OpenShift Container Platform on IBM Power® Virtual Server, you provide parameters to customize your cluster and the platform that hosts it. When you create the **install-config.yaml** file, you provide values for the required parameters through the command line. You can then modify the **install-config.yaml** file to customize your cluster further.

20.9.1. Available installation configuration parameters for IBM Power Virtual Server

The following tables specify the required, optional, and IBM Power Virtual Server-specific installation configuration parameters that you can set as part of the installation process.

**NOTE**

After installation, you cannot modify these parameters in the **install-config.yaml** file.

20.9.1.1. Required configuration parameters

Required installation configuration parameters are described in the following table:

Table 20.6. Required parameters

Parameter	Description	Values
-----------	-------------	--------

Parameter	Description	Values
apiVersion:	The API version for the install-config.yaml content. The current version is v1 . The installation program may also support older API versions.	String
baseDomain:	The base domain of your cloud provider. The base domain is used to create routes to your OpenShift Container Platform cluster components. The full DNS name for your cluster is a combination of the baseDomain and metadata.name parameter values that uses the <metadata.name> . <baseDomain> format.	A fully-qualified domain or subdomain name, such as example.com .
metadata:	Kubernetes resource ObjectMeta , from which only the name parameter is consumed.	Object
metadata: name:	The name of the cluster. DNS records for the cluster are all subdomains of {{.metadata.name}} . {{.baseDomain}} .	String of lowercase letters, hyphens (-), and periods (.), such as dev .
platform:	The configuration for the specific platform upon which to perform the installation: aws, baremetal, azure, gcp, ibmcloud, nutanix, openstack, powervs, vsphere , or {} . For additional information about platform . <platform> parameters, consult the table for your specific platform that follows.	Object

Parameter	Description	Values
<code>pullSecret:</code>	Get a pull secret from Red Hat OpenShift Cluster Manager to authenticate downloading container images for OpenShift Container Platform components from services such as Quay.io.	<pre>{ "auths":{ "cloud.openshift.com":{ "auth":"b3Blb=", "email":"you@example.com" }, "quay.io":{ "auth":"b3Blb=", "email":"you@example.com" } } }</pre>
<code>platform: powervs: userID:</code>	The UserID is the login for the user's IBM Cloud® account.	String. For example, existing_user_id .
<code>platform: powervs: powervsResourceGroup:</code>	The PowerVSResourceGroup is the resource group in which IBM Power® Virtual Server resources are created. If using an existing VPC, the existing VPC and subnets should be in this resource group.	String. For example, existing_resource_group .
<code>platform: powervs: region:</code>	Specifies the IBM Cloud® colo region where the cluster will be created.	String. For example, existing_region .
<code>platform: powervs: zone:</code>	Specifies the IBM Cloud® colo region where the cluster will be created.	String. For example, existing_zone .

20.9.1.2. Network configuration parameters


You can customize your installation configuration based on the requirements of your existing network infrastructure. For example, you can expand the IP address block for the cluster network or provide different IP address blocks than the defaults.


Only IPv4 addresses are supported.

**NOTE**

Globalnet is not supported with Red Hat OpenShift Data Foundation disaster recovery solutions. For regional disaster recovery scenarios, ensure that you use a nonoverlapping range of private IP addresses for the cluster and service networks in each cluster.

Table 20.7. Network parameters

Parameter	Description	Values
<code>networking:</code>	The configuration for the cluster network.	Object  NOTE You cannot modify parameters specified by the networking object after installation.
<code>networking: networkType:</code>	The Red Hat OpenShift Networking network plugin to install.	The default value is OVNKubernetes .
<code>networking: clusterNetwork:</code>	The IP address blocks for pods. The default value is 10.128.0.0/14 with a host prefix of /23 . If you specify multiple IP address blocks, the blocks must not overlap.	An array of objects. For example: <code>networking: clusterNetwork: - cidr: 10.128.0.0/14 hostPrefix: 23</code>
<code>networking: clusterNetwork: cidr:</code>	Required if you use networking.clusterNetwork . An IP address block. An IPv4 network.	An IP address block in Classless Inter-Domain Routing (CIDR) notation. The prefix length for an IPv4 block is between 0 and 32 .
<code>networking: clusterNetwork: hostPrefix:</code>	The subnet prefix length to assign to each individual node. For example, if hostPrefix is set to 23 then each node is assigned a /23 subnet out of the given cidr . A hostPrefix value of 23 provides 510 ($2^{(32 - 23)} - 2$) pod IP addresses.	A subnet prefix. The default value is 23 .
<code>networking: serviceNetwork:</code>	The IP address block for services. The default value is 172.30.0.0/16 . The OVN-Kubernetes network plugins supports only a single IP address block for the service network.	An array with an IP address block in CIDR format. For example: <code>networking: serviceNetwork: - 172.30.0.0/16</code>

Parameter	Description	Values
<code>networking: machineNetwork:</code>	The IP address blocks for machines.	An array of objects. For example: <code>networking: machineNetwork: - cidr: 10.0.0/16</code>
<code>networking: machineNetwork: cidr:</code>	Required if you use networking.machineNetwork . An IP address block. The default value is 10.0.0/16 for all platforms other than libvirt and IBM Power® Virtual Server. For libvirt, the default value is 192.168.126.0/24 . For IBM Power® Virtual Server, the default value is 192.168.0.0/24 .	An IP network block in CIDR notation. For example, 192.168.0.0/24 .  NOTE Set the networking.machineNetwork to match the CIDR that the preferred NIC resides in.


20.9.1.3. Optional configuration parameters


Optional installation configuration parameters are described in the following table:


Table 20.8. Optional parameters

Parameter	Description	Values
<code>additionalTrustBundle:</code>	A PEM-encoded X.509 certificate bundle that is added to the nodes' trusted certificate store. This trust bundle may also be used when a proxy has been configured.	String
<code>capabilities:</code>	Controls the installation of optional core cluster components. You can reduce the footprint of your OpenShift Container Platform cluster by disabling optional components. For more information, see the "Cluster capabilities" page in <i>Installing</i> .	String array
<code>capabilities: baselineCapabilitySet:</code>	Selects an initial set of optional capabilities to enable. Valid values are None , v4.11 , v4.12 and vCurrent . The default value is vCurrent .	String

Parameter	Description	Values
capabilities: additionalEnabledCapabilities:	Extends the set of optional capabilities beyond what you specify in baselineCapabilitySet . You may specify multiple capabilities in this parameter.	String array
cpuPartitioningMode:	Enables workload partitioning, which isolates OpenShift Container Platform services, cluster management workloads, and infrastructure pods to run on a reserved set of CPUs. Workload partitioning can only be enabled during installation and cannot be disabled after installation. While this field enables workload partitioning, it does not configure workloads to use specific CPUs. For more information, see the <i>Workload partitioning</i> page in the <i>Scalability and Performance</i> section.	None or AllNodes . None is the default value.
compute:	The configuration for the machines that comprise the compute nodes.	Array of MachinePool objects.
compute: architecture:	Determines the instruction set architecture of the machines in the pool. Currently, heterogeneous clusters are not supported, so all pools must specify the same architecture. Valid values are ppc64le (the default).	String

Parameter	Description	Values
<code>compute: hyperthreading:</code>	<p>Whether to enable or disable simultaneous multithreading, or hyperthreading, on compute machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.</p>  <p>IMPORTANT</p> <p>If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.</p>	Enabled or Disabled
<code>compute: smtLevel:</code>	The SMTLevel specifies the level of SMT to set to the control plane and compute machines. Valid values are 1, 2, 3, 4, 5, 6, 7, 8, off , and on .	String
<code>compute: name:</code>	Required if you use compute . The name of the machine pool.	worker
<code>compute: platform:</code>	Required if you use compute . Use this parameter to specify the cloud provider to host the worker machines. This parameter value must match the controlPlane.platform parameter value. Example usage, compute.platform.powervs.system .	aws, azure, gcp, ibmcloud, nutanix, openstack, powervs, vsphere, or {}
<code>compute: replicas:</code>	The number of compute machines, which are also known as worker machines, to provision.	A positive integer greater than or equal to 2 . The default value is 3 .
<code>featureSet:</code>	Enables the cluster for a feature set. A feature set is a collection of OpenShift Container Platform features that are not enabled by default. For more information about enabling a feature set during installation, see "Enabling features using feature gates".	String. The name of the feature set to enable, such as TechPreviewNoUpgrade .

Parameter	Description	Values
controlPlane:	The configuration for the machines that comprise the control plane.	Array of MachinePool objects.
controlPlane: architecture:	Determines the instruction set architecture of the machines in the pool. Currently, heterogeneous clusters are not supported, so all pools must specify the same architecture. Valid values are ppc64le (the default).	String
controlPlane: hyperthreading:	<p>Whether to enable or disable simultaneous multithreading, or hyperthreading, on control plane machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.</p> <div style="display: flex; align-items: center;">  <div> <p>IMPORTANT</p> <p>If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.</p> </div> </div>	Enabled or Disabled
controlPlane: name:	Required if you use controlPlane . The name of the machine pool.	master
controlPlane: platform:	Required if you use controlPlane . Use this parameter to specify the cloud provider that hosts the control plane machines. This parameter value must match the compute.platform parameter value. Example usage, controlPlane.platform.powersprocessors .	aws, azure, gcp, ibmcloud, nutanix, openstack, powersvs, vsphere , or {}
controlPlane: replicas:	The number of control plane machines to provision.	Supported values are 3 , or 1 when deploying single-node OpenShift.

Parameter	Description	Values
<code>credentialsMode:</code>	The Cloud Credential Operator (CCO) mode. If no mode is specified, the CCO dynamically tries to determine the capabilities of the provided credentials, with a preference for mint mode on the platforms where multiple modes are supported.	Mint, Passthrough, Manual or an empty string (<code>""</code>). ^[1]
<code>imageContentSources:</code>	Sources and repositories for the release-image content.	Array of objects. Includes a source and, optionally, mirrors , as described in the following rows of this table.
<code>imageContentSources: source:</code>	Required if you use imageContentSources . Specify the repository that users refer to, for example, in image pull specifications.	String
<code>imageContentSources: mirrors:</code>	Specify one or more repositories that may also contain the same images.	Array of strings
<code>publish:</code>	How to publish or expose the user-facing endpoints of your cluster, such as the Kubernetes API, OpenShift routes.	Internal or External . The default value is External . Setting this field to Internal is not supported on non-cloud platforms.
<code>sshKey:</code>	The SSH key to authenticate access to your cluster machines.  NOTE For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your ssh-agent process uses.	For example, sshKey: ssh-ed25519 AAAA...

Parameter	Description	Values
platform: powervs: vpcRegion:	Specifies the IBM Cloud® region in which to create VPC resources.	String. For example, existing_vpc_region .
platform: powervs: vpcSubnets:	Specifies existing subnets (by name) where cluster resources will be created.	String. For example, powervs_region_example_subnet .
platform: powervs: vpcName:	Specifies the IBM Cloud® name.	String. For example, existing_vpcName .
platform: powervs: serviceInstanceGUID:	The ServiceInstanceGUID is the ID of the Power IAAS instance created from the IBM Cloud® Catalog.	String. For example, existing_service_instance_GUID .
platform: powervs: clusterOSImage:	The ClusterOSImage is a pre-created IBM Power® Virtual Server boot image that overrides the default image for cluster nodes.	String. For example, existing_cluster_os_image .
platform: powervs: defaultMachinePlatform:	The DefaultMachinePlatform is the default configuration used when installing on IBM Power® Virtual Server for machine pools that do not define their own platform configuration.	String. For example, existing_machine_platform .
platform: powervs: memoryGiB:	The size of a virtual machine's memory, in GB.	The valid integer must be an integer number of GB that is at least 2 and no more than 64, depending on the machine type.
platform: powervs: procType:	The ProcType defines the processor sharing model for the instance.	The valid values are Capped, Dedicated, and Shared.

Parameter	Description	Values
<code>platform:</code> <code>powervs:</code> <code>processors:</code>	The Processors defines the processing units for the instance.	The number of processors must be from .5 to 32 cores. The processors must be in increments of .25.
<code>platform:</code> <code>powervs:</code> <code>sysType:</code>	The SysType defines the system type for the instance.	The system type must be either e980 or s922 .

1. Not all CCO modes are supported for all cloud providers. For more information about CCO modes, see the "Managing cloud provider credentials" entry in the *Authentication and authorization* content.

CHAPTER 21. INSTALLING ON OPENSTACK

21.1. PREPARING TO INSTALL ON OPENSTACK

You can install OpenShift Container Platform on Red Hat OpenStack Platform (RHOSP).

21.1.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).

21.1.2. Choosing a method to install OpenShift Container Platform on OpenStack

You can install OpenShift Container Platform on installer-provisioned or user-provisioned infrastructure. The default installation type uses installer-provisioned infrastructure, where the installation program provisions the underlying infrastructure for the cluster. You can also install OpenShift Container Platform on infrastructure that you provision. If you do not use infrastructure that the installation program provisions, you must manage and maintain the cluster resources yourself.

See [Installation process](#) for more information about installer-provisioned and user-provisioned installation processes.

21.1.2.1. Installing a cluster on installer-provisioned infrastructure

You can install a cluster on Red Hat OpenStack Platform (RHOSP) infrastructure that is provisioned by the OpenShift Container Platform installation program, by using one of the following methods:

- **Installing a cluster on OpenStack with customizations** You can install a customized cluster on RHOSP. The installation program allows for some customization to be applied at the installation stage. Many other customization options are available [post-installation](#).
- **Installing a cluster on OpenStack in a restricted network** You can install OpenShift Container Platform on RHOSP in a restricted or disconnected network by creating an internal mirror of the installation release content. You can use this method to install a cluster that does not require an active internet connection to obtain the software components. You can also use this installation method to ensure that your clusters only use container images that satisfy your organizational controls on external content.


21.1.2.2. Installing a cluster on user-provisioned infrastructure

You can install a cluster on RHOSP infrastructure that you provision, by using one of the following methods:

- **Installing a cluster on OpenStack on your own infrastructure** You can install OpenShift Container Platform on user-provisioned RHOSP infrastructure. By using this installation method, you can integrate your cluster with existing infrastructure and modifications. For installations on user-provisioned infrastructure, you must create all RHOSP resources, like Nova servers, Neutron ports, and security groups. You can use the provided Ansible playbooks to assist with the deployment process.

21.1.3. Scanning RHOSP endpoints for legacy HTTPS certificates

Beginning with OpenShift Container Platform 4.10, HTTPS certificates must contain subject alternative name (SAN) fields. Run the following script to scan each HTTPS endpoint in a Red Hat OpenStack Platform (RHOSP) catalog for legacy certificates that only contain the **CommonName** field.



IMPORTANT

OpenShift Container Platform does not check the underlying RHOSP infrastructure for legacy certificates prior to installation or updates. Use the provided script to check for these certificates yourself. Failing to update legacy certificates prior to installing or updating a cluster will result in cluster dysfunction.

Prerequisites

- On the machine where you run the script, have the following software:
 - Bash version 4.0 or greater
 - **grep**
 - [OpenStack client](#)
 - **jq**
 - [OpenSSL version 1.1.1l or greater](#)
- Populate the machine with RHOSP credentials for the target cloud.

Procedure

1. Save the following script to your machine:

```
#!/usr/bin/env bash

set -Eeuo pipefail

declare catalog san
catalog="$(mktemp)"
san="$(mktemp)"
readonly catalog san

declare invalid=0

openstack catalog list --format json --column Name --column Endpoints \
| jq -r '.[] | .Name as $name | .Endpoints[] | select(.interface=="public") | [$name, .interface, .url] | join(" ") \
| sort \
> "$catalog"

while read -r name interface url; do
# Ignore HTTP
if [[ ${url#"http://"} != "$url" ]]; then
continue
fi
```

```

# Remove the schema from the URL
noschema=${url#"https://"}

# If the schema was not HTTPS, error
if [[ "$noschema" == "$url" ]]; then
    echo "ERROR (unknown schema): $name $interface $url"
    exit 2
fi

# Remove the path and only keep host and port
noschema=${noschema%/*}
host=${noschema%:*}
port=${noschema##*:*}

# Add the port if was implicit
if [[ "$port" == "$host" ]]; then
    port='443'
fi

# Get the SAN fields
openssl s_client -showcerts -servername "$host" -connect "$host:$port" </dev/null
2>/dev/null \
| openssl x509 -noout -ext subjectAltName \
> "$san"

# openssl returns the empty string if no SAN is found.
# If a SAN is found, openssl is expected to return something like:
#
# X509v3 Subject Alternative Name:
#   DNS:standalone, DNS:osp1, IP Address:192.168.2.1, IP Address:10.254.1.2
if [[ "$(grep -c "Subject Alternative Name" "$san" || true)" -gt 0 ]]; then
    echo "PASS: $name $interface $url"
else
    invalid=$((invalid+1))
    echo "INVALID: $name $interface $url"
fi
done < "$catalog"

# clean up temporary files
rm "$catalog" "$san"

if [[ $invalid -gt 0 ]]; then
    echo "${invalid} legacy certificates were detected. Update your certificates to include a SAN
field."
    exit 1
else
    echo "All HTTPS certificates for this cloud are valid."
fi

```

2. Run the script.
3. Replace any certificates that the script reports as **INVALID** with certificates that contain SAN fields.



IMPORTANT

You must replace all legacy HTTPS certificates before you install OpenShift Container Platform 4.10 or update a cluster to that version. Legacy certificates will be rejected with the following message:

```
x509: certificate relies on legacy Common Name field, use SANs instead
```

21.1.3.1. Scanning RHOSP endpoints for legacy HTTPS certificates manually

Beginning with OpenShift Container Platform 4.10, HTTPS certificates must contain subject alternative name (SAN) fields. If you do not have access to the prerequisite tools that are listed in "Scanning RHOSP endpoints for legacy HTTPS certificates", perform the following steps to scan each HTTPS endpoint in a Red Hat OpenStack Platform (RHOSP) catalog for legacy certificates that only contain the **CommonName** field.



IMPORTANT

OpenShift Container Platform does not check the underlying RHOSP infrastructure for legacy certificates prior to installation or updates. Use the following steps to check for these certificates yourself. Failing to update legacy certificates prior to installing or updating a cluster will result in cluster dysfunction.

Procedure

1. On a command line, run the following command to view the URL of RHOSP public endpoints:

```
$ openstack catalog list
```

Record the URL for each HTTPS endpoint that the command returns.

2. For each public endpoint, note the host and the port.

TIP

Determine the host of an endpoint by removing the scheme, the port, and the path.

3. For each endpoint, run the following commands to extract the SAN field of the certificate:

- a. Set a **host** variable:

```
$ host=<host_name>
```

- b. Set a **port** variable:

```
$ port=<port_number>
```

If the URL of the endpoint does not have a port, use the value **443**.

- c. Retrieve the SAN field of the certificate:


```
$ openssl s_client -showcerts -servername "$host" -connect "$host:$port" </dev/null
2>/dev/null \
| openssl x509 -noout -ext subjectAltName
```

Example output

```
X509v3 Subject Alternative Name:
DNS:your.host.example.net
```

For each endpoint, look for output that resembles the previous example. If there is no output for an endpoint, the certificate of that endpoint is invalid and must be re-issued.



IMPORTANT

You must replace all legacy HTTPS certificates before you install OpenShift Container Platform 4.10 or update a cluster to that version. Legacy certificates are rejected with the following message:

```
x509: certificate relies on legacy Common Name field, use SANs instead
```

21.2. PREPARING TO INSTALL A CLUSTER THAT USES SR-IOV OR OVS-DPDK ON OPENSTACK

Before you install a OpenShift Container Platform cluster that uses single-root I/O virtualization (SR-IOV) or Open vSwitch with the Data Plane Development Kit (OVS-DPDK) on Red Hat OpenStack Platform (RHOSP), you must understand the requirements for each technology and then perform preparatory tasks.

21.2.1. Requirements for clusters on RHOSP that use either SR-IOV or OVS-DPDK

If you use SR-IOV or OVS-DPDK with your deployment, you must meet the following requirements:

- RHOSP compute nodes must use a flavor that supports huge pages.

21.2.1.1. Requirements for clusters on RHOSP that use SR-IOV

To use single-root I/O virtualization (SR-IOV) with your deployment, you must meet the following requirements:

- [Plan your Red Hat OpenStack Platform \(RHOSP\) SR-IOV deployment](#) .
- OpenShift Container Platform must support the NICs that you use. For a list of supported NICs, see "About Single Root I/O Virtualization (SR-IOV) hardware networks" in the "Hardware networks" subsection of the "Networking" documentation.
- For each node that will have an attached SR-IOV NIC, your RHOSP cluster must have:
 - One instance from the RHOSP quota
 - One port attached to the machines subnet
 - One port for each SR-IOV Virtual Function

- A flavor with at least 16 GB memory, 4 vCPUs, and 25 GB storage space
- SR-IOV deployments often employ performance optimizations, such as dedicated or isolated CPUs. For maximum performance, configure your underlying RHOSP deployment to use these optimizations, and then run OpenShift Container Platform compute machines on the optimized infrastructure.
- For more information about configuring performant RHOSP compute nodes, see [Configuring Compute nodes for performance](#).

21.2.1.2. Requirements for clusters on RHOSP that use OVS-DPDK

To use Open vSwitch with the Data Plane Development Kit (OVS-DPDK) with your deployment, you must meet the following requirements:

- Plan your Red Hat OpenStack Platform (RHOSP) OVS-DPDK deployment by referring to [Planning your OVS-DPDK deployment](#) in the Network Functions Virtualization Planning and Configuration Guide.
- Configure your RHOSP OVS-DPDK deployment according to [Configuring an OVS-DPDK deployment](#) in the Network Functions Virtualization Planning and Configuration Guide.

21.2.2. Preparing to install a cluster that uses SR-IOV

You must configure RHOSP before you install a cluster that uses SR-IOV on it.

When installing a cluster using SR-IOV, you must deploy clusters using cgroup v1. For more information, [Enabling Linux control group version 1 \(cgroup v1\)](#).



IMPORTANT

cgroup v1 is a deprecated feature. Deprecated functionality is still included in OpenShift Container Platform and continues to be supported; however, it will be removed in a future release of this product and is not recommended for new deployments.

For the most recent list of major functionality that has been deprecated or removed within OpenShift Container Platform, refer to the *Deprecated and removed features* section of the OpenShift Container Platform release notes.

21.2.2.1. Creating SR-IOV networks for compute machines

If your Red Hat OpenStack Platform (RHOSP) deployment supports [single root I/O virtualization \(SR-IOV\)](#), you can provision SR-IOV networks that compute machines run on.




NOTE

The following instructions entail creating an external flat network and an external, VLAN-based network that can be attached to a compute machine. Depending on your RHOSP deployment, other network types might be required.

Prerequisites

- Your cluster supports SR-IOV.

**NOTE**

If you are unsure about what your cluster supports, review the OpenShift Container Platform SR-IOV hardware networks documentation.

- You created radio and uplink provider networks as part of your RHOSP deployment. The names **radio** and **uplink** are used in all example commands to represent these networks.

Procedure

1. On a command line, create a radio RHOSP network:

```
$ openstack network create radio --provider-physical-network radio --provider-network-type flat --external
```

2. Create an uplink RHOSP network:

```
$ openstack network create uplink --provider-physical-network uplink --provider-network-type vlan --external
```

3. Create a subnet for the radio network:

```
$ openstack subnet create --network radio --subnet-range <radio_network_subnet_range> radio
```

4. Create a subnet for the uplink network:

```
$ openstack subnet create --network uplink --subnet-range <uplink_network_subnet_range> uplink
```

21.2.3. Preparing to install a cluster that uses OVS-DPDK

You must configure RHOSP before you install a cluster that uses SR-IOV on it.

- Complete [Creating a flavor and deploying an instance for OVS-DPDK](#) before you install a cluster on RHOSP.

After you perform preinstallation tasks, install your cluster by following the most relevant OpenShift Container Platform on RHOSP installation instructions. Then, perform the tasks under "Next steps" on this page.

21.2.4. Next steps

- For either type of deployment:
 - [Configure the Node Tuning Operator with huge pages support](#) .
- To complete SR-IOV configuration after you deploy your cluster:
 - [Install the SR-IOV Operator](#) .
 - [Configure your SR-IOV network device](#) .
 - [Create SR-IOV compute machines](#) .

- Consult the following references after you deploy your cluster to improve its performance:
 - [A test pod template for clusters that use OVS-DPDK on OpenStack](#).
 - [A test pod template for clusters that use SR-IOV on OpenStack](#).
 - [A performance profile template for clusters that use OVS-DPDK on OpenStack](#).

21.3. INSTALLING A CLUSTER ON OPENSTACK WITH CUSTOMIZATIONS

In OpenShift Container Platform version 4.16, you can install a customized cluster on Red Hat OpenStack Platform (RHOSP). To customize the installation, modify parameters in the **install-config.yaml** before you install the cluster.

21.3.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You verified that OpenShift Container Platform 4.16 is compatible with your RHOSP version by using the [Supported platforms for OpenShift clusters](#) section. You can also compare platform support across different versions by viewing the [OpenShift Container Platform on RHOSP support matrix](#).
- You have a storage service installed in RHOSP, such as block storage (Cinder) or object storage (Swift). Object storage is the recommended storage technology for OpenShift Container Platform registry cluster deployment. For more information, see [Optimizing storage](#).
- You understand performance and scalability practices for cluster scaling, control plane sizing, and etcd. For more information, see [Recommended practices for scaling the cluster](#).
- You have the metadata service enabled in RHOSP.

21.3.2. Resource guidelines for installing OpenShift Container Platform on RHOSP


To support an OpenShift Container Platform installation, your Red Hat OpenStack Platform (RHOSP) quota must meet the following requirements:

Table 21.1. Recommended resources for a default OpenShift Container Platform cluster on RHOSP

Resource	Value
Floating IP addresses	3
Ports	15
Routers	1
Subnets	1

Resource	Value
RAM	88 GB
vCPUs	22
Volume storage	275 GB
Instances	7
Security groups	3
Security group rules	60
Server groups	2 - plus 1 for each additional availability zone in each machine pool

A cluster might function with fewer than recommended resources, but its performance is not guaranteed.



IMPORTANT

If RHOSP object storage (Swift) is available and operated by a user account with the **swiftoperator** role, it is used as the default backend for the OpenShift Container Platform image registry. In this case, the volume storage requirement is 175 GB. Swift space requirements vary depending on the size of the image registry.



NOTE

By default, your security group and security group rule quotas might be low. If you encounter problems, run **openstack quota set --secgroups 3 --secgroup-rules 60 <project>** as an administrator to increase them.

An OpenShift Container Platform deployment comprises control plane machines, compute machines, and a bootstrap machine.

21.3.2.1. Control plane machines

By default, the OpenShift Container Platform installation process creates three control plane machines.

Each machine requires:

- An instance from the RHOSP quota
- A port from the RHOSP quota
- A flavor with at least 16 GB memory and 4 vCPUs
- At least 100 GB storage space from the RHOSP quota

21.3.2.2. Compute machines

By default, the OpenShift Container Platform installation process creates three compute machines.

Each machine requires:

- An instance from the RHOSP quota
- A port from the RHOSP quota
- A flavor with at least 8 GB memory and 2 vCPUs
- At least 100 GB storage space from the RHOSP quota

TIP

Compute machines host the applications that you run on OpenShift Container Platform; aim to run as many as you can.

21.3.2.3. Bootstrap machine

During installation, a bootstrap machine is temporarily provisioned to stand up the control plane. After the production control plane is ready, the bootstrap machine is deprovisioned.

The bootstrap machine requires:

- An instance from the RHOSP quota
- A port from the RHOSP quota
- A flavor with at least 16 GB memory and 4 vCPUs
- At least 100 GB storage space from the RHOSP quota

21.3.2.4. Load balancing requirements for user-provisioned infrastructure

Before you install OpenShift Container Platform, you can provision your own API and application ingress load balancing infrastructure to use in place of the default, internal load balancing solution. In production scenarios, you can deploy the API and application Ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.




NOTE

If you want to deploy the API and application Ingress load balancers with a Red Hat Enterprise Linux (RHEL) instance, you must purchase the RHEL subscription separately.

The load balancing infrastructure must meet the following requirements:

1. **API load balancer.** Provides a common endpoint for users, both human and machine, to interact with and configure the platform. Configure the following conditions:
 - Layer 4 load balancing only. This can be referred to as Raw TCP or SSL Passthrough mode.
 - A stateless load balancing algorithm. The options vary based on the load balancer implementation.



IMPORTANT

Do not configure session persistence for an API load balancer. Configuring session persistence for a Kubernetes API server might cause performance issues from excess application traffic for your OpenShift Container Platform cluster and the Kubernetes API that runs inside the cluster.

Configure the following ports on both the front and back of the load balancers:

Table 21.2. API load balancer

Port	Back-end machines (pool members)	Internal	External	Description
6443	Bootstrap and control plane. You remove the bootstrap machine from the load balancer after the bootstrap machine initializes the cluster control plane. You must configure the /readyz endpoint for the API server health check probe.	X	X	Kubernetes API server
22623	Bootstrap and control plane. You remove the bootstrap machine from the load balancer after the bootstrap machine initializes the cluster control plane.	X		Machine config server



NOTE

The load balancer must be configured to take a maximum of 30 seconds from the time the API server turns off the **/readyz** endpoint to the removal of the API server instance from the pool. Within the time frame after **/readyz** returns an error or becomes healthy, the endpoint must have been removed or added. Probing every 5 or 10 seconds, with two successful requests to become healthy and three to become unhealthy, are well-tested values.

2. **Application Ingress load balancer.** Provides an ingress point for application traffic flowing in from outside the cluster. A working configuration for the Ingress router is required for an OpenShift Container Platform cluster.

Configure the following conditions:

- Layer 4 load balancing only. This can be referred to as Raw TCP or SSL Passthrough mode.
- A connection-based or session-based persistence is recommended, based on the options available and types of applications that will be hosted on the platform.

TIP

If the true IP address of the client can be seen by the application Ingress load balancer, enabling source IP-based session persistence can improve performance for applications that use end-to-end TLS encryption.

Configure the following ports on both the front and back of the load balancers:

Table 21.3. Application Ingress load balancer

Port	Back-end machines (pool members)	Internal	External	Description
443	The machines that run the Ingress Controller pods, compute, or worker, by default.	X	X	HTTPS traffic
80	The machines that run the Ingress Controller pods, compute, or worker, by default.	X	X	HTTP traffic

**NOTE**

If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application Ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes.

21.3.2.4.1. Example load balancer configuration for clusters that are deployed with user-managed load balancers

This section provides an example API and application Ingress load balancer configuration that meets the load balancing requirements for clusters that are deployed with user-managed load balancers. The sample is an `/etc/haproxy/haproxy.cfg` configuration for an HAProxy load balancer. The example is not meant to provide advice for choosing one load balancing solution over another.

In the example, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

**NOTE**

If you are using HAProxy as a load balancer and SELinux is set to **enforcing**, you must ensure that the HAProxy service can bind to the configured TCP port by running **setsebool -P haproxy_connect_any=1**.

Example 21.1. Sample API and application Ingress load balancer configuration

```
global
  log      127.0.0.1 local2
  pidfile  /var/run/haproxy.pid
  maxconn  4000
  daemon
defaults
  mode          http
  log           global
  option        dontlognull
  option http-server-close
  option        redispatch
  retries       3
  timeout http-request 10s
  timeout queue 1m
```



```

timeout connect    10s
timeout client     1m
timeout server     1m
timeout http-keep-alive 10s
timeout check      10s
maxconn           3000
listen api-server-6443 1
bind *:6443
mode tcp
option httpchk GET /readyz HTTP/1.0
option log-health-checks
balance roundrobin
server bootstrap bootstrap.ocp4.example.com:6443 verify none check check-ssl inter 10s fall 2
rise 3 backup 2
server master0 master0.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
fall 2 rise 3
server master1 master1.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
fall 2 rise 3
server master2 master2.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
fall 2 rise 3
listen machine-config-server-22623 3
bind *:22623
mode tcp
server bootstrap bootstrap.ocp4.example.com:22623 check inter 1s backup 4
server master0 master0.ocp4.example.com:22623 check inter 1s
server master1 master1.ocp4.example.com:22623 check inter 1s
server master2 master2.ocp4.example.com:22623 check inter 1s
listen ingress-router-443 5
bind *:443
mode tcp
balance source
server compute0 compute0.ocp4.example.com:443 check inter 1s
server compute1 compute1.ocp4.example.com:443 check inter 1s
listen ingress-router-80 6
bind *:80
mode tcp
balance source
server compute0 compute0.ocp4.example.com:80 check inter 1s
server compute1 compute1.ocp4.example.com:80 check inter 1s

```

- 1** Port **6443** handles the Kubernetes API traffic and points to the control plane machines.
- 2** **4** The bootstrap entries must be in place before the OpenShift Container Platform cluster installation and they must be removed after the bootstrap process is complete.
- 3** Port **22623** handles the machine config server traffic and points to the control plane machines.
- 5** Port **443** handles the HTTPS traffic and points to the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.
- 6** Port **80** handles the HTTP traffic and points to the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.

**NOTE**

If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application Ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes.

TIP

If you are using HAProxy as a load balancer, you can check that the **haproxy** process is listening on ports **6443**, **22623**, **443**, and **80** by running **netstat -nltp** on the HAProxy node.

21.3.3. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.

**IMPORTANT**

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

21.3.4. Enabling Swift on RHOSP

Swift is operated by a user account with the **swiftoperator** role. Add the role to an account before you run the installation program.

**IMPORTANT**

If [the Red Hat OpenStack Platform \(RHOSP\) object storage service](#), commonly known as Swift, is available, OpenShift Container Platform uses it as the image registry storage. If it is unavailable, the installation program relies on the RHOSP block storage service, commonly known as Cinder.

If Swift is present and you want to use it, you must enable access to it. If it is not present, or if you do not want to use it, skip this section.



IMPORTANT

RHOSP 17 sets the **rgw_max_attr_size** parameter of Ceph RGW to 256 characters. This setting causes issues with uploading container images to the OpenShift Container Platform registry. You must set the value of **rgw_max_attr_size** to at least 1024 characters.

Before installation, check if your RHOSP deployment is affected by this problem. If it is, reconfigure Ceph RGW.

Prerequisites

- You have a RHOSP administrator account on the target environment.
- The Swift service is installed.
- On [Ceph RGW](#), the **account in url** option is enabled.

Procedure

To enable Swift on RHOSP:

1. As an administrator in the RHOSP CLI, add the **swiftoperator** role to the account that will access Swift:

```
$ openstack role add --user <user> --project <project> swiftoperator
```

Your RHOSP deployment can now use Swift for the image registry.

21.3.5. Configuring an image registry with custom storage on clusters that run on RHOSP

After you install a cluster on Red Hat OpenStack Platform (RHOSP), you can use a Cinder volume that is in a specific availability zone for registry storage.

Procedure

1. Create a YAML file that specifies the storage class and availability zone to use. For example:

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: custom-csi-storageclass
provisioner: cinder.csi.openstack.org
volumeBindingMode: WaitForFirstConsumer
allowVolumeExpansion: true
parameters:
  availability: <availability_zone_name>
```



NOTE

OpenShift Container Platform does not verify the existence of the availability zone you choose. Verify the name of the availability zone before you apply the configuration.

- From a command line, apply the configuration:

```
$ oc apply -f <storage_class_file_name>
```

Example output

```
storageclass.storage.k8s.io/custom-csi-storageclass created
```

- Create a YAML file that specifies a persistent volume claim (PVC) that uses your storage class and the **openshift-image-registry** namespace. For example:

```
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: csi-pvc-imageregistry
  namespace: openshift-image-registry
  annotations:
    imageregistry.openshift.io: "true"
spec:
  accessModes:
  - ReadWriteOnce
  volumeMode: Filesystem
  resources:
    requests:
      storage: 100Gi
  storageClassName: <your_custom_storage_class>
```

- Enter the namespace **openshift-image-registry**. This namespace allows the Cluster Image Registry Operator to consume the PVC.
- Optional: Adjust the volume size.
- Enter the name of the storage class that you created.

- From a command line, apply the configuration:

```
$ oc apply -f <pvc_file_name>
```

Example output

```
persistentvolumeclaim/csi-pvc-imageregistry created
```

- Replace the original persistent volume claim in the image registry configuration with the new claim:

```
$ oc patch configs.imageregistry.operator.openshift.io/cluster --type 'json' -p='[{"op": "replace", "path": "/spec/storage/pvc/claim", "value": "csi-pvc-imageregistry"}]'
```

Example output

```
config.imageregistry.operator.openshift.io/cluster patched
```

Over the next several minutes, the configuration is updated.

Verification

To confirm that the registry is using the resources that you defined:

1. Verify that the PVC claim value is identical to the name that you provided in your PVC definition:

```
$ oc get configs.imageregistry.operator.openshift.io/cluster -o yaml
```

Example output

```
...
status:
  ...
  managementState: Managed
  pvc:
    claim: csi-pvc-imageregistry
  ...
```

2. Verify that the status of the PVC is **Bound**:

```
$ oc get pvc -n openshift-image-registry csi-pvc-imageregistry
```

Example output

```
NAME                STATUS  VOLUME                                     CAPACITY  ACCESS MODES
STORAGECLASS        AGE
csi-pvc-imageregistry Bound   pvc-72a8f9c9-f462-11e8-b6b6-fa163e18b7b5  100Gi
RWO                  custom-csi-storageclass  11m
```

21.3.6. Verifying external network access

The OpenShift Container Platform installation process requires external network access. You must provide an external network value to it, or deployment fails. Before you begin the process, verify that a network with the external router type exists in Red Hat OpenStack Platform (RHOSP).

Prerequisites

- [Configure OpenStack's networking service to have DHCP agents forward instances' DNS queries](#)

Procedure

1. Using the RHOSP CLI, verify the name and ID of the 'External' network:

```
$ openstack network list --long -c ID -c Name -c "Router Type"
```

Example output

```
+-----+-----+-----+
| ID                | Name          | Router Type |
+-----+-----+-----+
```

```
| 148a8023-62a7-4672-b018-003462f8d7dc | public_network | External |
+-----+-----+-----+
```

A network with an external router type appears in the network list. If at least one does not, see [Creating a default floating IP network](#) and [Creating a default provider network](#).

IMPORTANT

If the external network's CIDR range overlaps one of the default network ranges, you must change the matching network ranges in the **install-config.yaml** file before you start the installation process.

The default network ranges are:

Network	Range
machineNetwork	10.0.0.0/16
serviceNetwork	172.30.0.0/16
clusterNetwork	10.128.0.0/14



WARNING

If the installation program finds multiple networks with the same name, it sets one of them at random. To avoid this behavior, create unique names for resources in RHOSP.



NOTE

If the Neutron trunk service plugin is enabled, a trunk port is created by default. For more information, see [Neutron trunk port](#).

21.3.7. Defining parameters for the installation program

The OpenShift Container Platform installation program relies on a file that is called **clouds.yaml**. The file describes Red Hat OpenStack Platform (RHOSP) configuration parameters, including the project name, log in information, and authorization service URLs.

Procedure

1. Create the **clouds.yaml** file:
 - If your RHOSP distribution includes the Horizon web UI, generate a **clouds.yaml** file in it.



IMPORTANT

Remember to add a password to the **auth** field. You can also keep secrets in a [separate file](#) from **clouds.yaml**.

- If your RHOSP distribution does not include the Horizon web UI, or you do not want to use Horizon, create the file yourself. For detailed information about **clouds.yaml**, see [Config files](#) in the RHOSP documentation.

```
clouds:
  shiftstack:
    auth:
      auth_url: http://10.10.14.42:5000/v3
      project_name: shiftstack
      username: <username>
      password: <password>
      user_domain_name: Default
      project_domain_name: Default
  dev-env:
    region_name: RegionOne
    auth:
      username: <username>
      password: <password>
      project_name: 'devonly'
      auth_url: 'https://10.10.14.22:5001/v2.0'
```

2. If your RHOSP installation uses self-signed certificate authority (CA) certificates for endpoint authentication:
 - a. Copy the certificate authority file to your machine.
 - b. Add the **cacerts** key to the **clouds.yaml** file. The value must be an absolute, non-root-accessible path to the CA certificate:

```
clouds:
  shiftstack:
    ...
    cacert: "/etc/pki/ca-trust/source/anchors/ca.crt.pem"
```

TIP

After you run the installer with a custom CA certificate, you can update the certificate by editing the value of the **ca-cert.pem** key in the **cloud-provider-config** keymap. On a command line, run:

```
$ oc edit configmap -n openshift-config cloud-provider-config
```

3. Place the **clouds.yaml** file in one of the following locations:
 - a. The value of the **OS_CLIENT_CONFIG_FILE** environment variable
 - b. The current directory
 - c. A Unix-specific user configuration directory, for example **~/.config/openstack/clouds.yaml**

- d. A Unix-specific site configuration directory, for example `/etc/openstack/clouds.yaml`
The installation program searches for **clouds.yaml** in that order.

21.3.8. Setting OpenStack Cloud Controller Manager options

Optionally, you can edit the OpenStack Cloud Controller Manager (CCM) configuration for your cluster. This configuration controls how OpenShift Container Platform interacts with Red Hat OpenStack Platform (RHOSP).

For a complete list of configuration parameters, see the "OpenStack Cloud Controller Manager reference guide" page in the "Installing on OpenStack" documentation.

Procedure

1. If you have not already generated manifest files for your cluster, generate them by running the following command:

```
$ openshift-install --dir <destination_directory> create manifests
```

2. In a text editor, open the cloud-provider configuration manifest file. For example:


```
$ vi openshift/manifests/cloud-provider-config.yaml
```

3. Modify the options according to the CCM reference guide.
Configuring Octavia for load balancing is a common case. For example:

```
#...
[LoadBalancer]
lb-provider = "amphora" 1
floating-network-id="d3deb660-4190-40a3-91f1-37326fe6ec4a" 2
create-monitor = True 3
monitor-delay = 10s 4
monitor-timeout = 10s 5
monitor-max-retries = 1 6
#...
```

- 1 This property sets the Octavia provider that your load balancer uses. It accepts **"ovn"** or **"amphora"** as values. If you choose to use OVN, you must also set **lb-method** to **SOURCE_IP_PORT**.
- 2 This property is required if you want to use multiple external networks with your cluster. The cloud provider creates floating IP addresses on the network that is specified here.
- 3 This property controls whether the cloud provider creates health monitors for Octavia load balancers. Set the value to **True** to create health monitors. As of RHOSP 16.2, this feature is only available for the Amphora provider.
- 4 This property sets the frequency with which endpoints are monitored. The value must be in the **time.ParseDuration()** format. This property is required if the value of the **create-monitor** property is **True**.
- 5 This property sets the time that monitoring requests are open before timing out. The value must be in the **time.ParseDuration()** format. This property is required if the value of the **create-monitor** property is **True**.

- 6 This property defines how many successful monitoring requests are required before a load balancer is marked as online. The value must be an integer. This property is required if the



IMPORTANT

Prior to saving your changes, verify that the file is structured correctly. Clusters might fail if properties are not placed in the appropriate section.



IMPORTANT

You must set the value of the **create-monitor** property to **True** if you use services that have the value of the **.spec.externalTrafficPolicy** property set to **Local**. The OVN Octavia provider in RHOSP 16.2 does not support health monitors. Therefore, services that have **ETP** parameter values set to **Local** might not respond when the **lb-provider** value is set to **"ovn"**.

4. Save the changes to the file and proceed with installation.

TIP

You can update your cloud provider configuration after you run the installer. On a command line, run:

```
$ oc edit configmap -n openshift-config cloud-provider-config
```

After you save your changes, your cluster will take some time to reconfigure itself. The process is complete if none of your nodes have a **SchedulingDisabled** status.

21.3.9. Obtaining the installation program


Before you install OpenShift Container Platform, download the installation file on the host you are using for installation.

Prerequisites


- You have a computer that runs Linux or macOS, with 500 MB of local disk space.

Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select your infrastructure provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.

**IMPORTANT**

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.

**IMPORTANT**

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.

4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#) . This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

21.3.10. Creating the installation configuration file

You can customize the OpenShift Container Platform cluster you install on Red Hat OpenStack Platform (RHOSP).

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create the **install-config.yaml** file.
 - a. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install create install-config --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the directory name to store the files that the installation program creates.

When specifying the directory:

- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
- Use an empty directory. Some installation assets, such as bootstrap X.509 certificates, have short expiration intervals, therefore you must not reuse an installation directory. If

you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.


- b. At the prompts, provide the configuration details for your cloud:
 - i. Optional: Select an SSH key to use to access your cluster machines.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- ii. Select **openstack** as the platform to target.
 - iii. Specify the Red Hat OpenStack Platform (RHOSP) external network name to use for installing the cluster.
 - iv. Specify the floating IP address to use for external access to the OpenShift API.
 - v. Specify a RHOSP flavor with at least 16 GB RAM to use for control plane nodes and 8 GB RAM for compute nodes.
 - vi. Select the base domain to deploy the cluster to. All DNS records will be sub-domains of this base and will also include the cluster name.
 - vii. Enter a name for your cluster. The name must be 14 or fewer characters long.
2. Modify the **install-config.yaml** file. You can find more information about the available parameters in the "Installation configuration parameters" section.
 3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

Additional resources

- [Installation configuration parameters for OpenStack](#)

21.3.10.1. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of

them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure


1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.

**NOTE**

The installation program does not support the proxy **readinessEndpoints** field.


**NOTE**

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.

**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

21.3.10.2. Custom subnets in RHOSP deployments

Optionally, you can deploy a cluster on a Red Hat OpenStack Platform (RHOSP) subnet of your choice. The subnet's GUID is passed as the value of **platform.openstack.machinesSubnet** in the **install-config.yaml** file.

This subnet is used as the cluster's primary subnet. By default, nodes and ports are created on it. You can create nodes and ports on a different RHOSP subnet by setting the value of the **platform.openstack.machinesSubnet** property to the subnet's UUID.

Before you run the OpenShift Container Platform installer with a custom subnet, verify that your configuration meets the following requirements:


- The subnet that is used by **platform.openstack.machinesSubnet** has DHCP enabled.
- The CIDR of **platform.openstack.machinesSubnet** matches the CIDR of **networking.machineNetwork**.
- The installation program user has permission to create ports on this network, including ports with fixed IP addresses.

Clusters that use custom subnets have the following limitations:

- If you plan to install a cluster that uses floating IP addresses, the **platform.openstack.machinesSubnet** subnet must be attached to a router that is connected to the **externalNetwork** network.
- If the **platform.openstack.machinesSubnet** value is set in the **install-config.yaml** file, the installation program does not create a private network or subnet for your RHOSP machines.
- You cannot use the **platform.openstack.externalDNS** property at the same time as a custom subnet. To add DNS to a cluster that uses a custom subnet, configure DNS on the RHOSP network.

**NOTE**

By default, the API VIP takes x.x.x.5 and the Ingress VIP takes x.x.x.7 from your network's CIDR block. To override these default values, set values for **platform.openstack.apiVIPs** and **platform.openstack.ingressVIPs** that are outside of the DHCP allocation pool.

**IMPORTANT**

The CIDR ranges for networks are not adjustable after cluster installation. Red Hat does not provide direct guidance on determining the range during cluster installation because it requires careful consideration of the number of created pods per namespace.

21.3.10.3. Deploying a cluster with bare metal machines

If you want your cluster to use bare metal machines, modify the **install-config.yaml** file. Your cluster can have both control plane and compute machines running on bare metal, or just compute machines.

**NOTE**

Be sure that your **install-config.yaml** file reflects whether the RHOSP network that you use for bare metal workers supports floating IP addresses or not.

Prerequisites

- The RHOSP [Bare Metal service \(Ironic\)](#) is enabled and accessible via the RHOSP Compute API.
- Bare metal is available as [a RHOSP flavor](#).
- If your cluster runs on an RHOSP version that is more than 16.1.6 and less than 16.2.4, bare metal workers do not function due to a [known issue](#) that causes the metadata service to be unavailable for services on OpenShift Container Platform nodes.
- The RHOSP network supports both VM and bare metal server attachment.
- Your network configuration does not rely on a provider network. Provider networks are not supported.
- If you want to deploy the machines on a pre-existing network, a RHOSP subnet is provisioned.
- If you want to deploy the machines on an installer-provisioned network, the RHOSP Bare Metal service (Ironic) is able to listen for and interact with Preboot eXecution Environment (PXE) boot machines that run on tenant networks.
- You created an **install-config.yaml** file as part of the OpenShift Container Platform installation process.

Procedure

1. In the **install-config.yaml** file, edit the flavors for machines:
 - a. If you want to use bare-metal control plane machines, change the value of **controlPlane.platform.openstack.type** to a bare metal flavor.
 - b. Change the value of **compute.platform.openstack.type** to a bare metal flavor.

- c. If you want to deploy your machines on a pre-existing network, change the value of **platform.openstack.machinesSubnet** to the RHOSP subnet UUID of the network. Control plane and compute machines must use the same subnet.

An example bare metal `install-config.yaml` file

```
controlPlane:
  platform:
    openstack:
      type: <bare_metal_control_plane_flavor> 1
  ...

compute:
  - architecture: amd64
    hyperthreading: Enabled
    name: worker
    platform:
      openstack:
        type: <bare_metal_compute_flavor> 2
      replicas: 3
  ...

platform:
  openstack:
    machinesSubnet: <subnet_UUID> 3
  ...
```

- 1 If you want to have bare-metal control plane machines, change this value to a bare metal flavor.
- 2 Change this value to a bare metal flavor to use for compute machines.
- 3 If you want to use a pre-existing network, change this value to the UUID of the RHOSP subnet.

Use the updated **install-config.yaml** file to complete the installation process. The compute machines that are created during deployment use the flavor that you added to the file.



NOTE

The installer may time out while waiting for bare metal machines to boot.

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

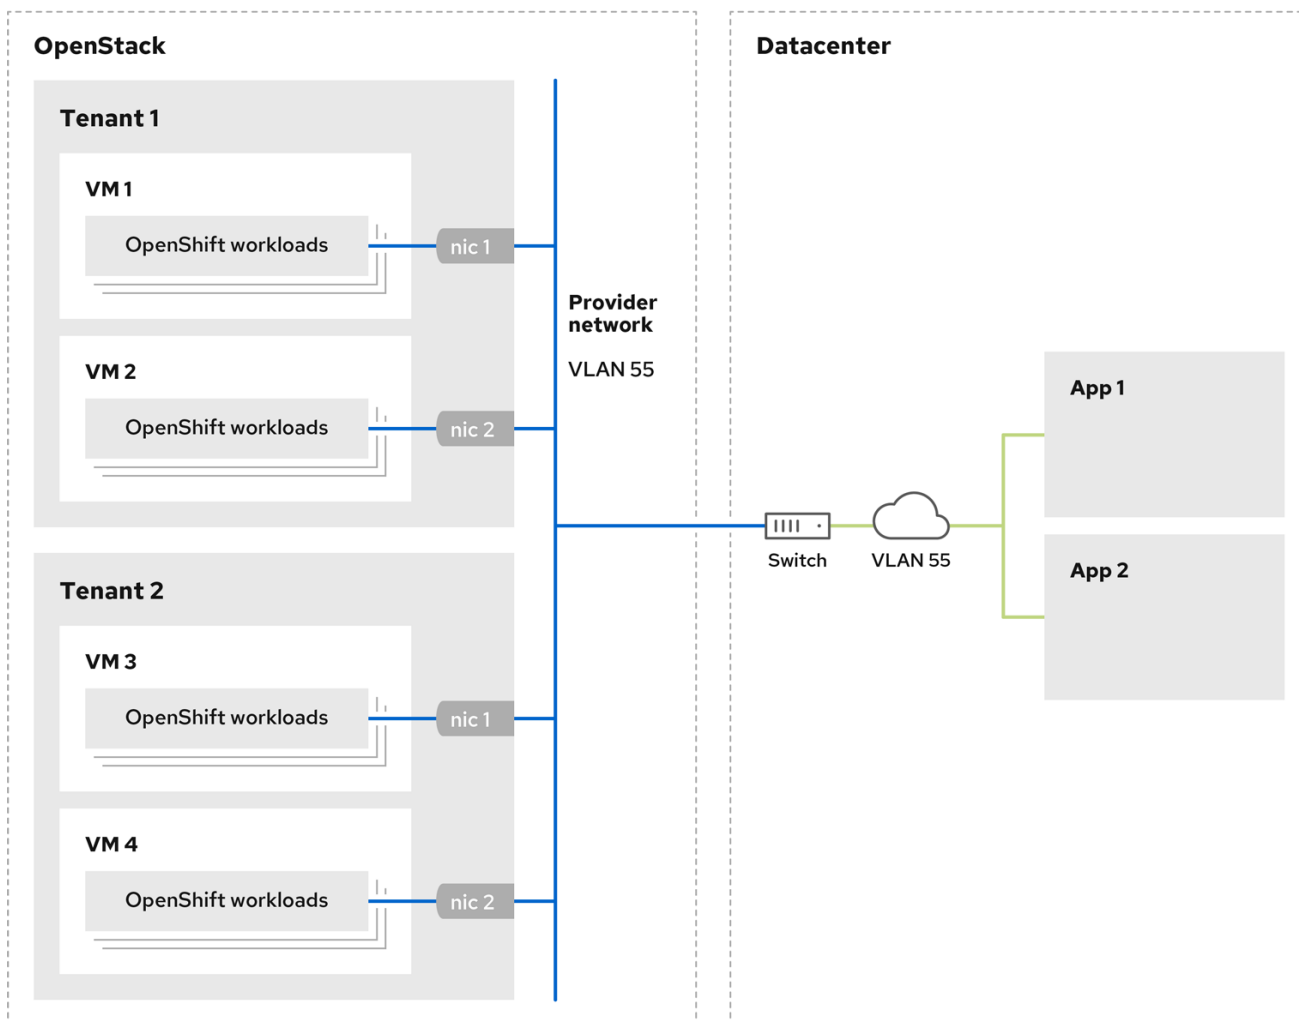
```
$ ./openshift-install wait-for install-complete --log-level debug
```

21.3.10.4. Cluster deployment on RHOSP provider networks

You can deploy your OpenShift Container Platform clusters on Red Hat OpenStack Platform (RHOSP) with a primary network interface on a provider network. Provider networks are commonly used to give projects direct access to a public network that can be used to reach the internet. You can also share provider networks among projects as part of the network creation process.

RHOSP provider networks map directly to an existing physical network in the data center. A RHOSP administrator must create them.

In the following example, OpenShift Container Platform workloads are connected to a data center by using a provider network:



170_OpenShift_0621

OpenShift Container Platform clusters that are installed on provider networks do not require tenant networks or floating IP addresses. The installer does not create these resources during installation.

Example provider network types include flat (untagged) and VLAN (802.1Q tagged).



NOTE

A cluster can support as many provider network connections as the network type allows. For example, VLAN networks typically support up to 4096 connections.

You can learn more about provider and tenant networks in [the RHOSP documentation](#).

21.3.10.4.1. RHOSP provider network requirements for cluster installation

Before you install an OpenShift Container Platform cluster, your Red Hat OpenStack Platform (RHOSP) deployment and provider network must meet a number of conditions:

- The [RHOSP networking service \(Neutron\) is enabled](#) and accessible through the RHOSP networking API.
- The RHOSP networking service has the [port security and allowed address pairs extensions enabled](#).
- The provider network can be shared with other tenants.

TIP

Use the **openstack network create** command with the **--share** flag to create a network that can be shared.

- The RHOSP project that you use to install the cluster must own the provider network, as well as an appropriate subnet.

TIP

To create a network for a project that is named "openshift," enter the following command


```
$ openstack network create --project openshift
```

To create a subnet for a project that is named "openshift," enter the following command

```
$ openstack subnet create --project openshift
```

To learn more about creating networks on RHOSP, read [the provider networks documentation](#).

If the cluster is owned by the **admin** user, you must run the installer as that user to create ports on the network.



IMPORTANT

Provider networks must be owned by the RHOSP project that is used to create the cluster. If they are not, the RHOSP Compute service (Nova) cannot request a port from that network.

- Verify that the provider network can reach the RHOSP metadata service IP address, which is **169.254.169.254** by default. Depending on your RHOSP SDN and networking service configuration, you might need to provide the route when you create the subnet. For example:

```
$ openstack subnet create --dhcp --host-route
destination=169.254.169.254/32,gateway=192.0.2.2 ...
```

- Optional: To secure the network, create [role-based access control \(RBAC\)](#) rules that limit network access to a single project.

21.3.10.4.2. Deploying a cluster that has a primary interface on a provider network

You can deploy an OpenShift Container Platform cluster that has its primary network interface on an Red Hat OpenStack Platform (RHOSP) provider network.

Prerequisites

- Your Red Hat OpenStack Platform (RHOSP) deployment is configured as described by "RHOSP provider network requirements for cluster installation".

Procedure

1. In a text editor, open the **install-config.yaml** file.
2. Set the value of the **platform.openstack.apiVIPs** property to the IP address for the API VIP.
3. Set the value of the **platform.openstack.ingressVIPs** property to the IP address for the Ingress VIP.
4. Set the value of the **platform.openstack.machinesSubnet** property to the UUID of the provider network subnet.
5. Set the value of the **networking.machineNetwork.cidr** property to the CIDR block of the provider network subnet.



IMPORTANT

The **platform.openstack.apiVIPs** and **platform.openstack.ingressVIPs** properties must both be unassigned IP addresses from the **networking.machineNetwork.cidr** block.

Section of an installation configuration file for a cluster that relies on a RHOSP provider network

```
...
platform:
  openstack:
    apiVIPs: 1
      - 192.0.2.13
    ingressVIPs: 2
      - 192.0.2.23
    machinesSubnet: fa806b2f-ac49-4bce-b9db-124bc64209bf
    # ...
  networking:
    machineNetwork:
      - cidr: 192.0.2.0/24
```

- 1 2 In OpenShift Container Platform 4.12 and later, the **apiVIP** and **ingressVIP** configuration settings are deprecated. Instead, use a list format to enter values in the **apiVIPs** and **ingressVIPs** configuration settings.

**WARNING**

You cannot set the **platform.openstack.externalNetwork** or **platform.openstack.externalDNS** parameters while using a provider network for the primary network interface.

When you deploy the cluster, the installer uses the **install-config.yaml** file to deploy the cluster on the provider network.

TIP

You can add additional networks, including provider networks, to the **platform.openstack.additionalNetworkIDs** list.

After you deploy your cluster, you can attach pods to additional networks. For more information, see [Understanding multiple networks](#).

21.3.10.5. Sample customized install-config.yaml file for RHOSP

The following example **install-config.yaml** files demonstrate all of the possible Red Hat OpenStack Platform (RHOSP) customization options.

**IMPORTANT**

This sample file is provided for reference only. You must obtain your **install-config.yaml** file by using the installation program.

Example 21.2. Example single stack install-config.yaml file

```

apiVersion: v1
baseDomain: example.com
controlPlane:
  name: master
  platform: {}
  replicas: 3
compute:
- name: worker
  platform:
    openstack:
      type: ml.large
  replicas: 3
metadata:
  name: example
networking:
  clusterNetwork:
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  machineNetwork:
  - cidr: 10.0.0.0/16

```

```

serviceNetwork:
- 172.30.0.0/16
networkType: OVNKubernetes
platform:
  openstack:
    cloud: mycloud
    externalNetwork: external
    computeFlavor: m1.xlarge
    apiFloatingIP: 128.0.0.1
fips: false
pullSecret: '{"auths": ...}'
sshKey: ssh-ed25519 AAAA...

```

Example 21.3. Example dual stack install-config.yaml file

```

apiVersion: v1
baseDomain: example.com
controlPlane:
  name: master
  platform: {}
  replicas: 3
compute:
- name: worker
  platform:
    openstack:
      type: ml.large
  replicas: 3
metadata:
  name: example
networking:
  clusterNetwork:
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  - cidr: fd01::/48
    hostPrefix: 64
  machineNetwork:
  - cidr: 192.168.25.0/24
  - cidr: fd2e:6f44:5dd8:c956::/64
  serviceNetwork:
  - 172.30.0.0/16
  - fd02::/112
  networkType: OVNKubernetes
platform:
  openstack:
    cloud: mycloud
    externalNetwork: external
    computeFlavor: m1.xlarge
  apiVIPs:
  - 192.168.25.10
  - fd2e:6f44:5dd8:c956:f816:3eff:fec3:5955
  ingressVIPs:
  - 192.168.25.132
  - fd2e:6f44:5dd8:c956:f816:3eff:fe40:aecb
  controlPlanePort:

```

```

fixedIPs:
- subnet:
  name: openshift-dual4
- subnet:
  name: openshift-dual6
network:
  name: openshift-dual
fips: false
pullSecret: '{"auths": ...}'
sshKey: ssh-ed25519 AAAA...

```

21.3.10.6. Optional: Configuring a cluster with dual-stack networking

You can create a dual-stack cluster on RHOSP. However, the dual-stack configuration is enabled only if you are using an RHOSP network with IPv4 and IPv6 subnets.



NOTE

RHOSP does not support the conversion of an IPv4 single-stack cluster to a dual-stack cluster network.

21.3.10.6.1. Deploying the dual-stack cluster

Procedure

1. Create a network with IPv4 and IPv6 subnets. The available address modes for the **ipv6-range-mode** and **ipv6-address-mode** fields are: **dhcpv6-stateful**, **dhcpv6-stateless**, and **slaac**.



NOTE

The dualstack network MTU must accommodate both the minimum MTU for IPv6, which is 1280, and the OVN-Kubernetes encapsulation overhead, which is 100.



NOTE

DHCP must be enabled on the subnets.

2. Create the API and Ingress VIPs ports.
3. Add the IPv6 subnet to the router to enable router advertisements. If you are using a provider network, you can enable router advertisements by adding the network as an external gateway, which also enables external connectivity.
4. To configure IPv4 and IPv6 address endpoints for cluster nodes, edit the **install-config.yaml** file. The following is an example of an **install-config.yaml** file:

Example install-config.yaml

```

apiVersion: v1
baseDomain: mydomain.test
compute:

```

```

- name: worker
  platform:
    openstack:
      type: m1.xlarge
  replicas: 3
controlPlane:
  name: master
  platform:
    openstack:
      type: m1.xlarge
  replicas: 3
metadata:
  name: mycluster
networking:
  machineNetwork: 1
  - cidr: "192.168.25.0/24"
  - cidr: "fd2e:6f44:5dd8:c956::/64"
  clusterNetwork: 2
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  - cidr: fd01::/48
    hostPrefix: 64
  serviceNetwork: 3
  - 172.30.0.0/16
  - fd02::/112
  platform:
    openstack:
      ingressVIPs: ['192.168.25.79', 'fd2e:6f44:5dd8:c956:f816:3eff:fef1:1bad'] 4
      apiVIPs: ['192.168.25.199', 'fd2e:6f44:5dd8:c956:f816:3eff:fe78:cf36'] 5
      controlPlanePort: 6
      fixedIPs: 7
      - subnet: 8
        name: subnet-v4
        id: subnet-v4-id
      - subnet: 9
        name: subnet-v6
        id: subnet-v6-id
      network: 10
        name: dualstack
        id: network-id

```

- 1 2 3 You must specify an IP address range for both the IPv4 and IPv6 address families.
- 4 Specify the virtual IP (VIP) address endpoints for the Ingress VIP services to provide an interface to the cluster.
- 5 Specify the virtual IP (VIP) address endpoints for the API VIP services to provide an interface to the cluster.
- 6 Specify the dual-stack network details that are used by all of the nodes across the cluster.
- 7 The CIDR of any subnet specified in this field must match the CIDRs listed on **networks.machineNetwork**.
- 8 9 You can specify a value for either **name** or **id**, or both.

- 10 Specifying the **network** under the **ControlPlanePort** field is optional.

Alternatively, if you want an IPv6 primary dual-stack cluster, edit the **install-config.yaml** file following the example below:

Example **install-config.yaml**

```

apiVersion: v1
baseDomain: mydomain.test
compute:
- name: worker
  platform:
    openstack:
      type: m1.xlarge
  replicas: 3
controlPlane:
  name: master
  platform:
    openstack:
      type: m1.xlarge
  replicas: 3
metadata:
  name: mycluster
networking:
  machineNetwork: 1
  - cidr: "fd2e:6f44:5dd8:c956::/64"
  - cidr: "192.168.25.0/24"
  clusterNetwork: 2
  - cidr: fd01::/48
    hostPrefix: 64
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  serviceNetwork: 3
  - fd02::/112
  - 172.30.0.0/16
platform:
  openstack:
    ingressVIPs: ["fd2e:6f44:5dd8:c956:f816:3eff:fef1:1bad", "192.168.25.79"] 4
    apiVIPs: ["fd2e:6f44:5dd8:c956:f816:3eff:fe78:cf36", "192.168.25.199"] 5
    controlPlanePort: 6
    fixedIPs: 7
    - subnet: 8
      name: subnet-v6
      id: subnet-v6-id
    - subnet: 9
      name: subnet-v4
      id: subnet-v4-id
    network: 10
      name: dualstack
      id: network-id

```

- 1 2 3 You must specify an IP address range for both the IPv4 and IPv6 address families.

- 4 Specify the virtual IP (VIP) address endpoints for the Ingress VIP services to provide an interface to the cluster.
- 5 Specify the virtual IP (VIP) address endpoints for the API VIP services to provide an interface to the cluster.
- 6 Specify the dual-stack network details that are used by all the nodes across the cluster.
- 7 The CIDR of any subnet specified in this field must match the CIDRs listed on **networks.machineNetwork**.
- 8 9 You can specify a value for either **name** or **id**, or both.
- 10 Specifying the **network** under the **ControlPlanePort** field is optional.

21.3.10.7. Installation configuration for a cluster on OpenStack with a user-managed load balancer

The following example **install-config.yaml** file demonstrates how to configure a cluster that uses an external, user-managed load balancer rather than the default internal load balancer.

```

apiVersion: v1
baseDomain: mydomain.test
compute:
- name: worker
  platform:
    openstack:
      type: m1.xlarge
      replicas: 3
controlPlane:
  name: master
  platform:
    openstack:
      type: m1.xlarge
      replicas: 3
metadata:
  name: mycluster
networking:
  clusterNetwork:
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  machineNetwork:
  - cidr: 192.168.10.0/24
platform:
  openstack:
    cloud: mycloud
    machinesSubnet: 8586bf1a-cc3c-4d40-bdf6-c243decc603a 1
    apiVIPs:
    - 192.168.10.5
    ingressVIPs:
    - 192.168.10.7
  loadBalancer:
    type: UserManaged 2

```


- 1 Regardless of which load balancer you use, the load balancer is deployed to this subnet.
- 2 The **UserManaged** value indicates that you are using an user-managed load balancer.

21.3.11. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.



NOTE

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.

**NOTE**

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the `ssh-agent` process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```

**NOTE**

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the `ssh-agent`:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

21.3.12. Enabling access to the environment

At deployment, all OpenShift Container Platform machines are created in a Red Hat OpenStack Platform (RHOSP)-tenant network. Therefore, they are not accessible directly in most RHOSP deployments.

You can configure OpenShift Container Platform API and application access by using floating IP addresses (FIPs) during installation. You can also complete an installation without configuring FIPs, but the installer will not configure a way to reach the API or applications externally.

21.3.12.1. Enabling access with floating IP addresses

Create floating IP (FIP) addresses for external access to the OpenShift Container Platform API and cluster applications.

Procedure

1. Using the Red Hat OpenStack Platform (RHOSP) CLI, create the API FIP:

```
$ openstack floating ip create --description "API <cluster_name>.<base_domain>"
<external_network>
```

2. Using the Red Hat OpenStack Platform (RHOSP) CLI, create the apps, or Ingress, FIP:

```
$ openstack floating ip create --description "Ingress <cluster_name>.<base_domain>"
<external_network>
```

3. Add records that follow these patterns to your DNS server for the API and Ingress FIPs:

```
api.<cluster_name>.<base_domain>. IN A <API_FIP>
*.apps.<cluster_name>.<base_domain>. IN A <apps_FIP>
```

NOTE

If you do not control the DNS server, you can access the cluster by adding the cluster domain names such as the following to your `/etc/hosts` file:

- `<api_floating_ip> api.<cluster_name>.<base_domain>`
- `<application_floating_ip> grafana-openshift-monitoring.apps.<cluster_name>.<base_domain>`
- `<application_floating_ip> prometheus-k8s-openshift-monitoring.apps.<cluster_name>.<base_domain>`
- `<application_floating_ip> oauth-openshift.apps.<cluster_name>.<base_domain>`
- `<application_floating_ip> console-openshift-console.apps.<cluster_name>.<base_domain>`
- `application_floating_ip integrated-oauth-server-openshift-authentication.apps.<cluster_name>.<base_domain>`

The cluster domain names in the `/etc/hosts` file grant access to the web console and the monitoring interface of your cluster locally. You can also use the `kubectl` or `oc`. You can access the user applications by using the additional entries pointing to the `<application_floating_ip>`. This action makes the API and applications accessible to only you, which is not suitable for production deployment, but does allow installation for development and testing.

4. Add the FIPs to the `install-config.yaml` file as the values of the following parameters:

- `platform.openstack.ingressFloatingIP`
- `platform.openstack.apiFloatingIP`

If you use these values, you must also enter an external network as the value of the **platform.openstack.externalNetwork** parameter in the **install-config.yaml** file.

TIP

You can make OpenShift Container Platform resources available outside of the cluster by assigning a floating IP address and updating your firewall configuration.

21.3.12.2. Completing installation without floating IP addresses

You can install OpenShift Container Platform on Red Hat OpenStack Platform (RHOSP) without providing floating IP addresses.

In the **install-config.yaml** file, do not define the following parameters:

- **platform.openstack.ingressFloatingIP**
- **platform.openstack.apiFloatingIP**

If you cannot provide an external network, you can also leave **platform.openstack.externalNetwork** blank. If you do not provide a value for **platform.openstack.externalNetwork**, a router is not created for you, and, without additional action, the installer will fail to retrieve an image from Glance. You must configure external connectivity on your own.

If you run the installer from a system that cannot reach the cluster API due to a lack of floating IP addresses or name resolution, installation fails. To prevent installation failure in these cases, you can use a proxy network or run the installer from a system that is on the same network as your machines.



NOTE

You can enable name resolution by creating DNS records for the API and Ingress ports. For example:

```
api.<cluster_name>.<base_domain>. IN A <api_port_IP>
*.apps.<cluster_name>.<base_domain>. IN A <ingress_port_IP>
```

If you do not control the DNS server, you can add the record to your **/etc/hosts** file. This action makes the API accessible to only you, which is not suitable for production deployment but does allow installation for development and testing.

21.3.13. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster.

- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.

Procedure

- Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

- 1 For **<installation_directory>**, specify the location of your customized **./install-config.yaml** file.
- 2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to **<installation_directory>/openshift_install.log**.




IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

21.3.14. Verifying cluster status

You can verify your OpenShift Container Platform cluster's status during or after installation.

Procedure

1. In the cluster environment, export the administrator's kubeconfig file:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server.

2. View the control plane and compute machines created after a deployment:

```
$ oc get nodes
```

3. View your cluster's version:

```
$ oc get clusterversion
```

4. View your Operators' status:

```
$ oc get clusteroperator
```

5. View all running pods in the cluster:

```
$ oc get pods -A
```

21.3.15. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container

Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

Additional resources

- See [Accessing the web console](#) for more details about accessing and understanding the OpenShift Container Platform web console.

21.3.16. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service

21.3.17. Next steps

- [Customize your cluster](#).
- If necessary, you can [opt out of remote health reporting](#) .
- If you need to enable external access to node ports, [configure ingress cluster traffic by using a node port](#).

- If you did not configure RHOSP to accept application traffic over floating IP addresses, [configure RHOSP access with floating IP addresses](#).

21.4. INSTALLING A CLUSTER ON OPENSTACK ON YOUR OWN INFRASTRUCTURE

In OpenShift Container Platform version 4.16, you can install a cluster on Red Hat OpenStack Platform (RHOSP) that runs on user-provisioned infrastructure.

Using your own infrastructure allows you to integrate your cluster with existing infrastructure and modifications. The process requires more labor on your part than installer-provisioned installations, because you must create all RHOSP resources, like Nova servers, Neutron ports, and security groups. However, Red Hat provides Ansible playbooks to help you in the deployment process.

21.4.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You verified that OpenShift Container Platform 4.16 is compatible with your RHOSP version by using the [Supported platforms for OpenShift clusters](#) section. You can also compare platform support across different versions by viewing the [OpenShift Container Platform on RHOSP support matrix](#).
- You have an RHOSP account where you want to install OpenShift Container Platform.
- You understand performance and scalability practices for cluster scaling, control plane sizing, and etcd. For more information, see [Recommended practices for scaling the cluster](#).
- On the machine from which you run the installation program, you have:
 - A single directory in which you can keep the files you create during the installation process
 - Python 3

21.4.2. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.



IMPORTANT

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.


21.4.3. Resource guidelines for installing OpenShift Container Platform on RHOSP

To support an OpenShift Container Platform installation, your Red Hat OpenStack Platform (RHOSP) quota must meet the following requirements:

Table 21.4. Recommended resources for a default OpenShift Container Platform cluster on RHOSP

Resource	Value
Floating IP addresses	3
Ports	15
Routers	1
Subnets	1
RAM	88 GB
vCPUs	22
Volume storage	275 GB
Instances	7
Security groups	3
Security group rules	60
Server groups	2 - plus 1 for each additional availability zone in each machine pool

A cluster might function with fewer than recommended resources, but its performance is not guaranteed.



IMPORTANT

If RHOSP object storage (Swift) is available and operated by a user account with the **swiftoperator** role, it is used as the default backend for the OpenShift Container Platform image registry. In this case, the volume storage requirement is 175 GB. Swift space requirements vary depending on the size of the image registry.

**NOTE**

By default, your security group and security group rule quotas might be low. If you encounter problems, run **`openstack quota set --secgroups 3 --secgroup-rules 60 <project>`** as an administrator to increase them.

An OpenShift Container Platform deployment comprises control plane machines, compute machines, and a bootstrap machine.

21.4.3.1. Control plane machines

By default, the OpenShift Container Platform installation process creates three control plane machines.

Each machine requires:

- An instance from the RHOSP quota
- A port from the RHOSP quota
- A flavor with at least 16 GB memory and 4 vCPUs
- At least 100 GB storage space from the RHOSP quota

21.4.3.2. Compute machines

By default, the OpenShift Container Platform installation process creates three compute machines.

Each machine requires:

- An instance from the RHOSP quota
- A port from the RHOSP quota
- A flavor with at least 8 GB memory and 2 vCPUs
- At least 100 GB storage space from the RHOSP quota

TIP

Compute machines host the applications that you run on OpenShift Container Platform; aim to run as many as you can.

21.4.3.3. Bootstrap machine

During installation, a bootstrap machine is temporarily provisioned to stand up the control plane. After the production control plane is ready, the bootstrap machine is deprovisioned.

The bootstrap machine requires:

- An instance from the RHOSP quota
- A port from the RHOSP quota
- A flavor with at least 16 GB memory and 4 vCPUs
- At least 100 GB storage space from the RHOSP quota

21.4.4. Downloading playbook dependencies

The Ansible playbooks that simplify the installation process on user-provisioned infrastructure require several Python modules. On the machine where you will run the installer, add the modules' repositories and then download them.



NOTE

These instructions assume that you are using Red Hat Enterprise Linux (RHEL) 8.

Prerequisites

- Python 3 is installed on your machine.

Procedure

1. On a command line, add the repositories:

- a. Register with Red Hat Subscription Manager:

```
$ sudo subscription-manager register # If not done already
```

- b. Pull the latest subscription data:

```
$ sudo subscription-manager attach --pool=$YOUR_POOLID # If not done already
```

- c. Disable the current repositories:

```
$ sudo subscription-manager repos --disable=* # If not done already
```

- d. Add the required repositories:

```
$ sudo subscription-manager repos \
  --enable=rhel-8-for-x86_64-baseos-rpms \
  --enable=openstack-16-tools-for-rhel-8-x86_64-rpms \
  --enable=ansible-2.9-for-rhel-8-x86_64-rpms \
  --enable=rhel-8-for-x86_64-appstream-rpms
```

2. Install the modules:

```
$ sudo yum install python3-openstackclient ansible python3-openstacksdk python3-netaddr
ansible-collections-openstack
```

3. Ensure that the **python** command points to **python3**:

```
$ sudo alternatives --set python /usr/bin/python3
```

21.4.5. Downloading the installation playbooks

Download Ansible playbooks that you can use to install OpenShift Container Platform on your own Red Hat OpenStack Platform (RHOSP) infrastructure.

Prerequisites

- The curl command-line tool is available on your machine.

Procedure

- To download the playbooks to your working directory, run the following script from a command line:

```
$ xargs -n 1 curl -O <<< '  
  https://raw.githubusercontent.com/openshift/installer/release-  
4.16/upi/openstack/bootstrap.yaml  
  https://raw.githubusercontent.com/openshift/installer/release-  
4.16/upi/openstack/common.yaml  
  https://raw.githubusercontent.com/openshift/installer/release-  
4.16/upi/openstack/compute-nodes.yaml  
  https://raw.githubusercontent.com/openshift/installer/release-4.16/upi/openstack/control-  
plane.yaml  
  https://raw.githubusercontent.com/openshift/installer/release-4.16/upi/openstack/down-  
bootstrap.yaml  
  https://raw.githubusercontent.com/openshift/installer/release-4.16/upi/openstack/down-  
compute-nodes.yaml  
  https://raw.githubusercontent.com/openshift/installer/release-4.16/upi/openstack/down-  
control-plane.yaml  
  https://raw.githubusercontent.com/openshift/installer/release-4.16/upi/openstack/down-  
network.yaml  
  https://raw.githubusercontent.com/openshift/installer/release-4.16/upi/openstack/down-  
security-groups.yaml  
  https://raw.githubusercontent.com/openshift/installer/release-4.16/upi/openstack/down-  
containers.yaml  
  https://raw.githubusercontent.com/openshift/installer/release-  
4.16/upi/openstack/inventory.yaml  
  https://raw.githubusercontent.com/openshift/installer/release-  
4.16/upi/openstack/network.yaml  
  https://raw.githubusercontent.com/openshift/installer/release-  
4.16/upi/openstack/security-groups.yaml  
  https://raw.githubusercontent.com/openshift/installer/release-4.16/upi/openstack/update-  
network-resources.yaml'
```

The playbooks are downloaded to your machine.



IMPORTANT

During the installation process, you can modify the playbooks to configure your deployment.

Retain all playbooks for the life of your cluster. You must have the playbooks to remove your OpenShift Container Platform cluster from RHOSP.



IMPORTANT

You must match any edits you make in the **bootstrap.yaml**, **compute-nodes.yaml**, **control-plane.yaml**, **network.yaml**, and **security-groups.yaml** files to the corresponding playbooks that are prefixed with **down-**. For example, edits to the **bootstrap.yaml** file must be reflected in the **down-bootstrap.yaml** file, too. If you do not edit both files, the supported cluster removal process will fail.

21.4.6. Obtaining the installation program

Before you install OpenShift Container Platform, download the installation file on the host you are using for installation.

Prerequisites

- You have a computer that runs Linux or macOS, with 500 MB of local disk space.


Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select your infrastructure provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.



IMPORTANT

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.



IMPORTANT

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.

4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#). This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

21.4.7. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.



NOTE

You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.



NOTE

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.

**NOTE**

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```

**NOTE**

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

21.4.8. Creating the Red Hat Enterprise Linux CoreOS (RHCOS) image

The OpenShift Container Platform installation program requires that a Red Hat Enterprise Linux CoreOS (RHCOS) image be present in the Red Hat OpenStack Platform (RHOSP) cluster. Retrieve the latest RHCOS image, then upload it using the RHOSP CLI.


Prerequisites

- The RHOSP CLI is installed.

Procedure

1. Log in to the Red Hat Customer Portal's [Product Downloads page](#).

- Under **Version**, select the most recent release of OpenShift Container Platform 4.16 for Red Hat Enterprise Linux (RHEL) 8.



IMPORTANT

The RHCOS images might not change with every release of OpenShift Container Platform. You must download images with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Use the image versions that match your OpenShift Container Platform version if they are available.

- Download the *Red Hat Enterprise Linux CoreOS (RHCOS) - OpenStack Image (QCOW)* .
- Decompress the image.




NOTE

You must decompress the RHOSP image before the cluster can use it. The name of the downloaded file might not contain a compression extension, like **.gz** or **.tgz**. To find out if or how the file is compressed, in a command line, enter:

```
$ file <name_of_downloaded_file>
```

- From the image that you downloaded, create an image that is named **rhcos** in your cluster by using the RHOSP CLI:

```
$ openstack image create --container-format=bare --disk-format=qcow2 --file rhcos-
${RHCOS_VERSION}-openstack.qcow2 rhcos
```



IMPORTANT

Depending on your RHOSP environment, you might be able to upload the image in either **.raw** or **.qcow2** formats. If you use Ceph, you must use the **.raw** format.



WARNING

If the installation program finds multiple images with the same name, it chooses one of them at random. To avoid this behavior, create unique names for resources in RHOSP.

After you upload the image to RHOSP, it is usable in the installation process.

21.4.9. Verifying external network access

The OpenShift Container Platform installation process requires external network access. You must provide an external network value to it, or deployment fails. Before you begin the process, verify that a network with the external router type exists in Red Hat OpenStack Platform (RHOSP).

Prerequisites

- [Configure OpenStack's networking service to have DHCP agents forward instances' DNS queries](#)

Procedure

1. Using the RHOSP CLI, verify the name and ID of the 'External' network:

```
$ openstack network list --long -c ID -c Name -c "Router Type"
```

Example output

```
+-----+-----+-----+
| ID              | Name          | Router Type |
+-----+-----+-----+
| 148a8023-62a7-4672-b018-003462f8d7dc | public_network | External    |
+-----+-----+-----+
```

A network with an external router type appears in the network list. If at least one does not, see [Creating a default floating IP network](#) and [Creating a default provider network](#).



NOTE

If the Neutron trunk service plugin is enabled, a trunk port is created by default. For more information, see [Neutron trunk port](#).

21.4.10. Enabling access to the environment

At deployment, all OpenShift Container Platform machines are created in a Red Hat OpenStack Platform (RHOSP)-tenant network. Therefore, they are not accessible directly in most RHOSP deployments.

You can configure OpenShift Container Platform API and application access by using floating IP addresses (FIPs) during installation. You can also complete an installation without configuring FIPs, but the installer will not configure a way to reach the API or applications externally.

21.4.10.1. Enabling access with floating IP addresses

Create floating IP (FIP) addresses for external access to the OpenShift Container Platform API, cluster applications, and the bootstrap process.

Procedure

1. Using the Red Hat OpenStack Platform (RHOSP) CLI, create the API FIP:

```
$ openstack floating ip create --description "API <cluster_name>.<base_domain>"
<external_network>
```

2. Using the Red Hat OpenStack Platform (RHOSP) CLI, create the apps, or Ingress, FIP:

```
$ openstack floating ip create --description "Ingress <cluster_name>.<base_domain>"
<external_network>
```

- By using the Red Hat OpenStack Platform (RHOSP) CLI, create the bootstrap FIP:

```
$ openstack floating ip create --description "bootstrap machine" <external_network>
```

- Add records that follow these patterns to your DNS server for the API and Ingress FIPs:

```
api.<cluster_name>.<base_domain>. IN A <API_FIP>
*.apps.<cluster_name>.<base_domain>. IN A <apps_FIP>
```

NOTE

If you do not control the DNS server, you can access the cluster by adding the cluster domain names such as the following to your `/etc/hosts` file:

- `<api_floating_ip> api.<cluster_name>.<base_domain>`
- `<application_floating_ip> grafana-openshift-monitoring.apps.<cluster_name>.<base_domain>`
- `<application_floating_ip> prometheus-k8s-openshift-monitoring.apps.<cluster_name>.<base_domain>`
- `<application_floating_ip> oauth-openshift.apps.<cluster_name>.<base_domain>`
- `<application_floating_ip> console-openshift-console.apps.<cluster_name>.<base_domain>`
- `application_floating_ip integrated-oauth-server-openshift-authentication.apps.<cluster_name>.<base_domain>`

The cluster domain names in the `/etc/hosts` file grant access to the web console and the monitoring interface of your cluster locally. You can also use the `kubectl` or `oc`. You can access the user applications by using the additional entries pointing to the `<application_floating_ip>`. This action makes the API and applications accessible to only you, which is not suitable for production deployment, but does allow installation for development and testing.

- Add the FIPs to the `inventory.yaml` file as the values of the following variables:

- `os_api_fip`
- `os_bootstrap_fip`
- `os_ingress_fip`

If you use these values, you must also enter an external network as the value of the `os_external_network` variable in the `inventory.yaml` file.

TIP

You can make OpenShift Container Platform resources available outside of the cluster by assigning a floating IP address and updating your firewall configuration.

21.4.10.2. Completing installation without floating IP addresses

You can install OpenShift Container Platform on Red Hat OpenStack Platform (RHOSP) without providing floating IP addresses.

In the **inventory.yaml** file, do not define the following variables:

- **os_api_fip**
- **os_bootstrap_fip**
- **os_ingress_fip**

If you cannot provide an external network, you can also leave **os_external_network** blank. If you do not provide a value for **os_external_network**, a router is not created for you, and, without additional action, the installer will fail to retrieve an image from Glance. Later in the installation process, when you create network resources, you must configure external connectivity on your own.

If you run the installer with the **wait-for** command from a system that cannot reach the cluster API due to a lack of floating IP addresses or name resolution, installation fails. To prevent installation failure in these cases, you can use a proxy network or run the installer from a system that is on the same network as your machines.



NOTE

You can enable name resolution by creating DNS records for the API and Ingress ports. For example:

```
api.<cluster_name>.<base_domain>. IN A <api_port_IP>
*.apps.<cluster_name>.<base_domain>. IN A <ingress_port_IP>
```

If you do not control the DNS server, you can add the record to your **/etc/hosts** file. This action makes the API accessible to only you, which is not suitable for production deployment but does allow installation for development and testing.

21.4.11. Defining parameters for the installation program

The OpenShift Container Platform installation program relies on a file that is called **clouds.yaml**. The file describes Red Hat OpenStack Platform (RHOSP) configuration parameters, including the project name, log in information, and authorization service URLs.

Procedure

1. Create the **clouds.yaml** file:
 - If your RHOSP distribution includes the Horizon web UI, generate a **clouds.yaml** file in it.



IMPORTANT

Remember to add a password to the **auth** field. You can also keep secrets in [a separate file](#) from **clouds.yaml**.

- If your RHOSP distribution does not include the Horizon web UI, or you do not want to use Horizon, create the file yourself. For detailed information about **clouds.yaml**, see [Config files](#) in the RHOSP documentation.

```

clouds:
  shiftstack:
    auth:
      auth_url: http://10.10.14.42:5000/v3
      project_name: shiftstack
      username: <username>
      password: <password>
      user_domain_name: Default
      project_domain_name: Default
    dev-env:
      region_name: RegionOne
      auth:
        username: <username>
        password: <password>
        project_name: 'devonly'
        auth_url: 'https://10.10.14.22:5001/v2.0'

```

2. If your RHOSP installation uses self-signed certificate authority (CA) certificates for endpoint authentication:
 - a. Copy the certificate authority file to your machine.
 - b. Add the **cacerts** key to the **clouds.yaml** file. The value must be an absolute, non-root-accessible path to the CA certificate:

```

clouds:
  shiftstack:
    ...
    cacert: "/etc/pki/ca-trust/source/anchors/ca.crt.pem"

```

TIP

After you run the installer with a custom CA certificate, you can update the certificate by editing the value of the **ca-cert.pem** key in the **cloud-provider-config** keymap. On a command line, run:

```
$ oc edit configmap -n openshift-config cloud-provider-config
```

3. Place the **clouds.yaml** file in one of the following locations:
 - a. The value of the **OS_CLIENT_CONFIG_FILE** environment variable
 - b. The current directory
 - c. A Unix-specific user configuration directory, for example **~/.config/openstack/clouds.yaml**
 - d. A Unix-specific site configuration directory, for example **/etc/openstack/clouds.yaml**

The installation program searches for **clouds.yaml** in that order.

21.4.12. Creating network resources on RHOSP

Create the network resources that an OpenShift Container Platform on Red Hat OpenStack Platform (RHOSP) installation on your own infrastructure requires. To save time, run supplied Ansible playbooks that generate security groups, networks, subnets, routers, and ports.

Prerequisites

- You downloaded the modules in "Downloading playbook dependencies".
- You downloaded the playbooks in "Downloading the installation playbooks".

Procedure

1. For a dual stack cluster deployment, edit the **inventory.yaml** file and uncomment the **os_subnet6** attribute.
2. On a command line, create the network resources by running the following command:

```
$ ansible-playbook -i inventory.yaml network.yaml
```



NOTE

The API and Ingress VIP fields will be overwritten in the **inventory.yaml** playbook with the IP addresses assigned to the network ports.



NOTE

The resources created by the **network.yaml** playbook are deleted by the **down-network.yaml** playbook.

21.4.13. Creating the installation configuration file

You can customize the OpenShift Container Platform cluster you install on Red Hat OpenStack Platform (RHOSP).

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create the **install-config.yaml** file.
 - a. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install create install-config --dir <installation_directory> 1
```



For **<installation_directory>**, specify the directory name to store the files that the installation program creates.

When specifying the directory:

- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
- Use an empty directory. Some installation assets, such as bootstrap X.509 certificates,

have short expiration intervals, therefore you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.


- b. At the prompts, provide the configuration details for your cloud:
 - i. Optional: Select an SSH key to use to access your cluster machines.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- ii. Select **openstack** as the platform to target.
 - iii. Specify the Red Hat OpenStack Platform (RHOSP) external network name to use for installing the cluster.
 - iv. Specify the floating IP address to use for external access to the OpenShift API.
 - v. Specify a RHOSP flavor with at least 16 GB RAM to use for control plane nodes and 8 GB RAM for compute nodes.
 - vi. Select the base domain to deploy the cluster to. All DNS records will be sub-domains of this base and will also include the cluster name.
 - vii. Enter a name for your cluster. The name must be 14 or fewer characters long.
2. Modify the **install-config.yaml** file. You can find more information about the available parameters in the "Installation configuration parameters" section.
 3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

You now have the file **install-config.yaml** in the directory that you specified.

Additional resources

- [Installation configuration parameters for OpenStack](#)

21.4.13.1. Custom subnets in RHOSP deployments

Optionally, you can deploy a cluster on a Red Hat OpenStack Platform (RHOSP) subnet of your choice. The subnet's GUID is passed as the value of **platform.openstack.machinesSubnet** in the **install-config.yaml** file.

This subnet is used as the cluster’s primary subnet. By default, nodes and ports are created on it. You can create nodes and ports on a different RHOSP subnet by setting the value of the **platform.openstack.machinesSubnet** property to the subnet’s UUID.

Before you run the OpenShift Container Platform installer with a custom subnet, verify that your configuration meets the following requirements:

- The subnet that is used by **platform.openstack.machinesSubnet** has DHCP enabled.
- The CIDR of **platform.openstack.machinesSubnet** matches the CIDR of **networking.machineNetwork**.
- The installation program user has permission to create ports on this network, including ports with fixed IP addresses.

Clusters that use custom subnets have the following limitations:

- If you plan to install a cluster that uses floating IP addresses, the **platform.openstack.machinesSubnet** subnet must be attached to a router that is connected to the **externalNetwork** network.
- If the **platform.openstack.machinesSubnet** value is set in the **install-config.yaml** file, the installation program does not create a private network or subnet for your RHOSP machines.
- You cannot use the **platform.openstack.externalDNS** property at the same time as a custom subnet. To add DNS to a cluster that uses a custom subnet, configure DNS on the RHOSP network.



NOTE

By default, the API VIP takes x.x.x.5 and the Ingress VIP takes x.x.x.7 from your network’s CIDR block. To override these default values, set values for **platform.openstack.apiVIPs** and **platform.openstack.ingressVIPs** that are outside of the DHCP allocation pool.




IMPORTANT

The CIDR ranges for networks are not adjustable after cluster installation. Red Hat does not provide direct guidance on determining the range during cluster installation because it requires careful consideration of the number of created pods per namespace.

21.4.13.2. Sample customized install-config.yaml file for RHOSP

The following example **install-config.yaml** files demonstrate all of the possible Red Hat OpenStack Platform (RHOSP) customization options.



IMPORTANT

This sample file is provided for reference only. You must obtain your **install-config.yaml** file by using the installation program.

Example 21.4. Example single stack install-config.yaml file

```
apiVersion: v1
baseDomain: example.com
controlPlane:
```

```
name: master
platform: {}
replicas: 3
compute:
- name: worker
  platform:
    openstack:
      type: ml.large
  replicas: 3
metadata:
  name: example
networking:
  clusterNetwork:
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  machineNetwork:
  - cidr: 10.0.0.0/16
  serviceNetwork:
  - 172.30.0.0/16
  networkType: OVNKubernetes
platform:
  openstack:
    cloud: mycloud
    externalNetwork: external
    computeFlavor: m1.xlarge
    apiFloatingIP: 128.0.0.1
fips: false
pullSecret: '{"auths": ...}'
sshKey: ssh-ed25519 AAAA...
```

Example 21.5. Example dual stack install-config.yaml file

```
apiVersion: v1
baseDomain: example.com
controlPlane:
  name: master
  platform: {}
  replicas: 3
compute:
- name: worker
  platform:
    openstack:
      type: ml.large
  replicas: 3
metadata:
  name: example
networking:
  clusterNetwork:
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  - cidr: fd01::/48
    hostPrefix: 64
  machineNetwork:
  - cidr: 192.168.25.0/24
```



```

- cidr: fd2e:6f44:5dd8:c956::/64
serviceNetwork:
- 172.30.0.0/16
- fd02::/112
networkType: OVNKubernetes
platform:
openstack:
  cloud: mycloud
  externalNetwork: external
  computeFlavor: m1.xlarge
  apiVIPs:
  - 192.168.25.10
  - fd2e:6f44:5dd8:c956:f816:3eff:fec3:5955
  ingressVIPs:
  - 192.168.25.132
  - fd2e:6f44:5dd8:c956:f816:3eff:fe40:aecb
  controlPlanePort:
  fixedIPs:
  - subnet:
    name: openshift-dual4
  - subnet:
    name: openshift-dual6
  network:
    name: openshift-dual
fips: false
pullSecret: '{"auths": ...}'
sshKey: ssh-ed25519 AAAA...

```

21.4.13.3. Setting a custom subnet for machines

The IP range that the installation program uses by default might not match the Neutron subnet that you create when you install OpenShift Container Platform. If necessary, update the CIDR value for new machines by editing the installation configuration file.

Prerequisites

- You have the **install-config.yaml** file that was generated by the OpenShift Container Platform installation program.
- You have Python 3 installed.

Procedure

1. On a command line, browse to the directory that contains the **install-config.yaml** and **inventory.yaml** files.
2. From that directory, either run a script to edit the **install-config.yaml** file or update the file manually:
 - To set the value by using a script, run the following command:

```

$ python -c 'import yaml
path = "install-config.yaml"
data = yaml.safe_load(open(path))

```

```

inventory = yaml.safe_load(open("inventory.yaml"))["all"]["hosts"]["localhost"]
machine_net = [{"cidr": inventory["os_subnet_range"]}
api_vips = [inventory["os_apiVIP"]]
ingress_vips = [inventory["os_ingressVIP"]]
ctrl_plane_port = {"network": {"name": inventory["os_network"]}, "fixedIPs": [{"subnet":
{"name": inventory["os_subnet"]}]}
if inventory.get("os_subnet6"): ❶
    machine_net.append({"cidr": inventory["os_subnet6_range"]})
    api_vips.append(inventory["os_apiVIP6"])
    ingress_vips.append(inventory["os_ingressVIP6"])
    data["networking"]["networkType"] = "OVNKubernetes"
    data["networking"]["clusterNetwork"].append({"cidr": inventory["cluster_network6_cidr"],
"hostPrefix": inventory["cluster_network6_prefix"]})
    data["networking"]["serviceNetwork"].append(inventory["service_subnet6_range"])
    ctrl_plane_port["fixedIPs"].append({"subnet": {"name": inventory["os_subnet6"]})}
data["networking"]["machineNetwork"] = machine_net
data["platform"]["openstack"]["apiVIPs"] = api_vips
data["platform"]["openstack"]["ingressVIPs"] = ingress_vips
data["platform"]["openstack"]["controlPlanePort"] = ctrl_plane_port
del data["platform"]["openstack"]["externalDNS"]
open(path, "w").write(yaml.dump(data, default_flow_style=False))

```

- ❶ Applies to dual stack (IPv4/IPv6) environments.

21.4.13.4. Emptying compute machine pools

To proceed with an installation that uses your own infrastructure, set the number of compute machines in the installation configuration file to zero. Later, you create these machines manually.

Prerequisites

- You have the **install-config.yaml** file that was generated by the OpenShift Container Platform installation program.

Procedure

- On a command line, browse to the directory that contains **install-config.yaml**.
- From that directory, either run a script to edit the **install-config.yaml** file or update the file manually:
 - To set the value by using a script, run:

```

$ python -c '
import yaml;
path = "install-config.yaml";
data = yaml.safe_load(open(path));
data["compute"][0]["replicas"] = 0;
open(path, "w").write(yaml.dump(data, default_flow_style=False))'

```

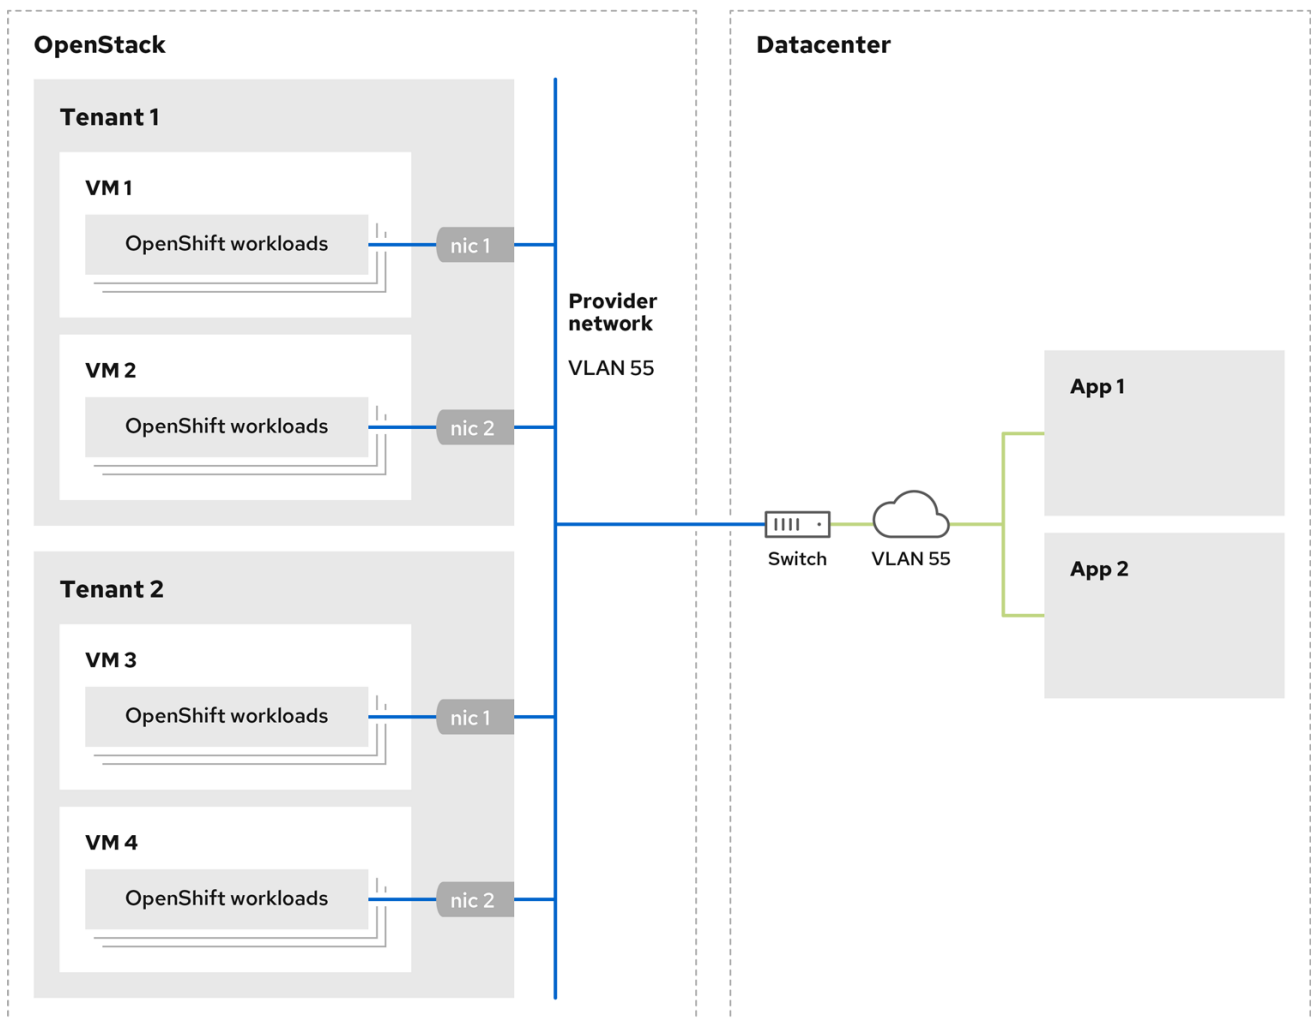
- To set the value manually, open the file and set the value of **compute.<first entry>.replicas** to **0**.

21.4.13.5. Cluster deployment on RHOSP provider networks

You can deploy your OpenShift Container Platform clusters on Red Hat OpenStack Platform (RHOSP) with a primary network interface on a provider network. Provider networks are commonly used to give projects direct access to a public network that can be used to reach the internet. You can also share provider networks among projects as part of the network creation process.

RHOSP provider networks map directly to an existing physical network in the data center. A RHOSP administrator must create them.

In the following example, OpenShift Container Platform workloads are connected to a data center by using a provider network:



170_OpenShift_0621

OpenShift Container Platform clusters that are installed on provider networks do not require tenant networks or floating IP addresses. The installer does not create these resources during installation.

Example provider network types include flat (untagged) and VLAN (802.1Q tagged).



NOTE

A cluster can support as many provider network connections as the network type allows. For example, VLAN networks typically support up to 4096 connections.

You can learn more about provider and tenant networks in [the RHOSP documentation](#).

21.4.13.5.1. RHOSP provider network requirements for cluster installation

Before you install an OpenShift Container Platform cluster, your Red Hat OpenStack Platform (RHOSP) deployment and provider network must meet a number of conditions:

- The [RHOSP networking service \(Neutron\) is enabled](#) and accessible through the RHOSP networking API.
- The RHOSP networking service has the [port security and allowed address pairs extensions enabled](#).
- The provider network can be shared with other tenants.

TIP

Use the **openstack network create** command with the **--share** flag to create a network that can be shared.

- The RHOSP project that you use to install the cluster must own the provider network, as well as an appropriate subnet.

TIP

To create a network for a project that is named "openshift," enter the following command


```
$ openstack network create --project openshift
```

To create a subnet for a project that is named "openshift," enter the following command

```
$ openstack subnet create --project openshift
```

To learn more about creating networks on RHOSP, read [the provider networks documentation](#).

If the cluster is owned by the **admin** user, you must run the installer as that user to create ports on the network.



IMPORTANT

Provider networks must be owned by the RHOSP project that is used to create the cluster. If they are not, the RHOSP Compute service (Nova) cannot request a port from that network.

- Verify that the provider network can reach the RHOSP metadata service IP address, which is **169.254.169.254** by default. Depending on your RHOSP SDN and networking service configuration, you might need to provide the route when you create the subnet. For example:

```
$ openstack subnet create --dhcp --host-route
destination=169.254.169.254/32,gateway=192.0.2.2 ...
```

- Optional: To secure the network, create [role-based access control \(RBAC\)](#) rules that limit network access to a single project.

21.4.13.5.2. Deploying a cluster that has a primary interface on a provider network

You can deploy an OpenShift Container Platform cluster that has its primary network interface on an Red Hat OpenStack Platform (RHOSP) provider network.

Prerequisites

- Your Red Hat OpenStack Platform (RHOSP) deployment is configured as described by "RHOSP provider network requirements for cluster installation".

Procedure

1. In a text editor, open the **install-config.yaml** file.
2. Set the value of the **platform.openstack.apiVIPs** property to the IP address for the API VIP.
3. Set the value of the **platform.openstack.ingressVIPs** property to the IP address for the Ingress VIP.
4. Set the value of the **platform.openstack.machinesSubnet** property to the UUID of the provider network subnet.
5. Set the value of the **networking.machineNetwork.cidr** property to the CIDR block of the provider network subnet.



IMPORTANT

The **platform.openstack.apiVIPs** and **platform.openstack.ingressVIPs** properties must both be unassigned IP addresses from the **networking.machineNetwork.cidr** block.

Section of an installation configuration file for a cluster that relies on a RHOSP provider network

```
...
platform:
  openstack:
    apiVIPs: ①
      - 192.0.2.13
    ingressVIPs: ②
      - 192.0.2.23
    machinesSubnet: fa806b2f-ac49-4bce-b9db-124bc64209bf
    # ...
  networking:
    machineNetwork:
      - cidr: 192.0.2.0/24
```

- ① ② In OpenShift Container Platform 4.12 and later, the **apiVIP** and **ingressVIP** configuration settings are deprecated. Instead, use a list format to enter values in the **apiVIPs** and **ingressVIPs** configuration settings.

**WARNING**

You cannot set the **platform.openstack.externalNetwork** or **platform.openstack.externalDNS** parameters while using a provider network for the primary network interface.

When you deploy the cluster, the installer uses the **install-config.yaml** file to deploy the cluster on the provider network.

TIP

You can add additional networks, including provider networks, to the **platform.openstack.additionalNetworkIDs** list.

After you deploy your cluster, you can attach pods to additional networks. For more information, see [Understanding multiple networks](#).

21.4.14. Creating the Kubernetes manifest and Ignition config files

Because you must modify some cluster definition files and manually start the cluster machines, you must generate the Kubernetes manifest and Ignition config files that the cluster needs to configure the machines.

The installation configuration file transforms into the Kubernetes manifests. The manifests wrap into the Ignition configuration files, which are later used to configure the cluster machines.

**IMPORTANT**

- The Ignition config files that the OpenShift Container Platform installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

Prerequisites

- You obtained the OpenShift Container Platform installation program.
- You created the **install-config.yaml** installation configuration file.

Procedure

1. Change to the directory that contains the OpenShift Container Platform installation program and generate the Kubernetes manifests for the cluster:

```
$ ./openshift-install create manifests --dir <installation_directory> 1
```

- 1 For **<installation_directory>**, specify the installation directory that contains the **install-config.yaml** file you created.

2. Remove the Kubernetes manifest files that define the control plane machines, compute machine sets, and control plane machine sets:

```
$ rm -f openshift/99_openshift-cluster-api_master-machines-*.yaml openshift/99_openshift-cluster-api_worker-machineset-*.yaml openshift/99_openshift-machine-api_master-control-plane-machine-set.yaml
```

Because you create and manage these resources yourself, you do not have to initialize them.

- You can preserve the compute machine set files to create compute machines by using the machine API, but you must update references to them to match your environment.
3. Check that the **mastersSchedulable** parameter in the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** Kubernetes manifest file is set to **false**. This setting prevents pods from being scheduled on the control plane machines:
 - a. Open the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** file.
 - b. Locate the **mastersSchedulable** parameter and ensure that it is set to **false**.
 - c. Save and exit the file.
 4. To create the Ignition configuration files, run the following command from the directory that contains the installation program:

```
$ ./openshift-install create ignition-configs --dir <installation_directory> 1
```

- 1 For **<installation_directory>**, specify the same installation directory.

Ignition config files are created for the bootstrap, control plane, and compute nodes in the installation directory. The **kubeadmin-password** and **kubeconfig** files are created in the **./<installation_directory>/auth** directory:

```
.
├── auth
│   ├── kubeadmin-password
│   └── kubeconfig
├── bootstrap.ign
├── master.ign
├── metadata.json
└── worker.ign
```

5. Export the metadata file's **infraID** key as an environment variable:

```
$ export INFRA_ID=$(jq -r .infraID metadata.json)
```

TIP

Extract the **infraID** key from **metadata.json** and use it as a prefix for all of the RHOSP resources that you create. By doing so, you avoid name conflicts when making multiple deployments in the same project.

21.4.15. Preparing the bootstrap Ignition files

The OpenShift Container Platform installation process relies on bootstrap machines that are created from a bootstrap Ignition configuration file.

Edit the file and upload it. Then, create a secondary bootstrap Ignition configuration file that Red Hat OpenStack Platform (RHOSP) uses to download the primary file.

Prerequisites

- You have the bootstrap Ignition file that the installer program generates, **bootstrap.ign**.
- The infrastructure ID from the installer's metadata file is set as an environment variable (**\$INFRA_ID**).
 - If the variable is not set, see **Creating the Kubernetes manifest and Ignition config files**
- You have an HTTP(S)-accessible way to store the bootstrap Ignition file.
 - The documented procedure uses the RHOSP image service (Glance), but you can also use the RHOSP storage service (Swift), Amazon S3, an internal HTTP server, or an ad hoc Nova server.

Procedure

1. Run the following Python script. The script modifies the bootstrap Ignition file to set the hostname and, if available, CA certificate file when it runs:

```
import base64
import json
import os

with open('bootstrap.ign', 'r') as f:
    ignition = json.load(f)

files = ignition['storage'].get('files', [])

infra_id = os.environ.get('INFRA_ID', 'openshift').encode()
hostname_b64 = base64.standard_b64encode(infra_id + b'-bootstrap\n').decode().strip()
files.append(
    {
        'path': '/etc/hostname',
        'mode': 420,
        'contents': {
            'source': 'data:text/plain;charset=utf-8;base64,' + hostname_b64
        }
    }
)
```



```

ca_cert_path = os.environ.get('OS_CACERT', "")
if ca_cert_path:
    with open(ca_cert_path, 'r') as f:
        ca_cert = f.read().encode()
        ca_cert_b64 = base64.standard_b64encode(ca_cert).decode().strip()

    files.append(
        {
            'path': '/opt/openshift/tls/cloud-ca-cert.pem',
            'mode': 420,
            'contents': {
                'source': 'data:text/plain;charset=utf-8;base64,' + ca_cert_b64
            }
        }
    )

ignition['storage']['files'] = files;

with open('bootstrap.ign', 'w') as f:
    json.dump(ignition, f)

```

- Using the RHOSP CLI, create an image that uses the bootstrap Ignition file:

```

$ openstack image create --disk-format=raw --container-format=bare --file bootstrap.ign
<image_name>

```


- Get the image's details:

```

$ openstack image show <image_name>

```

Make a note of the **file** value; it follows the pattern **v2/images/<image_ID>/file**.



NOTE

Verify that the image you created is active.

- Retrieve the image service's public address:

```

$ openstack catalog show image

```

- Combine the public address with the image **file** value and save the result as the storage location. The location follows the pattern **<image_service_public_URL>/v2/images/<image_ID>/file**.

- Generate an auth token and save the token ID:

```

$ openstack token issue -c id -f value

```

- Insert the following content into a file called **\$INFRA_ID-bootstrap-ignition.json** and edit the placeholders to match your own values:

```

{
  "ignition": {

```

```

"config": {
  "merge": [{
    "source": "<storage_url>", ❶
    "httpHeaders": [{
      "name": "X-Auth-Token", ❷
      "value": "<token_ID>" ❸
    }]
  }],
  "security": {
    "tls": {
      "certificateAuthorities": [{
        "source": "data:text/plain;charset=utf-8;base64,<base64_encoded_certificate>" ❹
      }]
    }
  },
  "version": "3.2.0"
}

```

- ❶ Replace the value of **ignition.config.merge.source** with the bootstrap Ignition file storage URL.
- ❷ Set **name** in **httpHeaders** to **"X-Auth-Token"**.
- ❸ Set **value** in **httpHeaders** to your token's ID.
- ❹ If the bootstrap Ignition file server uses a self-signed certificate, include the base64-encoded certificate.

8. Save the secondary Ignition config file.

The bootstrap Ignition data will be passed to RHOSP during installation.



WARNING

The bootstrap Ignition file contains sensitive information, like **clouds.yaml** credentials. Ensure that you store it in a secure place, and delete it after you complete the installation process.

21.4.16. Creating control plane Ignition config files on RHOSP

Installing OpenShift Container Platform on Red Hat OpenStack Platform (RHOSP) on your own infrastructure requires control plane Ignition config files. You must create multiple config files.



NOTE

As with the bootstrap Ignition configuration, you must explicitly define a hostname for each control plane machine.

Prerequisites

- The infrastructure ID from the installation program's metadata file is set as an environment variable (**\$INFRA_ID**).
 - If the variable is not set, see "Creating the Kubernetes manifest and Ignition config files".

Procedure

- On a command line, run the following Python script:

```
$ for index in $(seq 0 2); do
  MASTER_HOSTNAME="$INFRA_ID-master-$index\n"
  python -c "import base64, json, sys;
  ignition = json.load(sys.stdin);
  storage = ignition.get('storage', {});
  files = storage.get('files', []);
  files.append({'path': '/etc/hostname', 'mode': 420, 'contents': {'source':
  'data:text/plain;charset=utf-8;base64,' +
  base64.standard_b64encode(b'$MASTER_HOSTNAME').decode().strip(), 'verification': {}},
  'filesystem': 'root'});
  storage['files'] = files;
  ignition['storage'] = storage
  json.dump(ignition, sys.stdout) <master.ign >"$INFRA_ID-master-$index-ignition.json"
done
```

You now have three control plane Ignition files: **<INFRA_ID>-master-0-ignition.json**, **<INFRA_ID>-master-1-ignition.json**, and **<INFRA_ID>-master-2-ignition.json**.

21.4.17. Updating network resources on RHOSP

Update the network resources that an OpenShift Container Platform on Red Hat OpenStack Platform (RHOSP) installation on your own infrastructure requires.

Prerequisites

- Python 3 is installed on your machine.
- You downloaded the modules in "Downloading playbook dependencies".
- You downloaded the playbooks in "Downloading the installation playbooks".

Procedure

1. Optional: Add an external network value to the **inventory.yaml** playbook:

Example external network value in the **inventory.yaml** Ansible Playbook

```
...
# The public network providing connectivity to the cluster. If not
# provided, the cluster external connectivity must be provided in another
# way.
```

```
# Required for os_api_fip, os_ingress_fip, os_bootstrap_fip.
os_external_network: 'external'
```

...



IMPORTANT

If you did not provide a value for **os_external_network** in the **inventory.yaml** file, you must ensure that VMs can access Glance and an external connection yourself.

- Optional: Add external network and floating IP (FIP) address values to the **inventory.yaml** playbook:

Example FIP values in the **inventory.yaml** Ansible Playbook

...

```
# OpenShift API floating IP address. If this value is non-empty, the
# corresponding floating IP will be attached to the Control Plane to
# serve the OpenShift API.
os_api_fip: '203.0.113.23'

# OpenShift Ingress floating IP address. If this value is non-empty, the
# corresponding floating IP will be attached to the worker nodes to serve
# the applications.
os_ingress_fip: '203.0.113.19'

# If this value is non-empty, the corresponding floating IP will be
# attached to the bootstrap machine. This is needed for collecting logs
# in case of install failure.
os_bootstrap_fip: '203.0.113.20'
```



IMPORTANT

If you do not define values for **os_api_fip** and **os_ingress_fip**, you must perform postinstallation network configuration.

If you do not define a value for **os_bootstrap_fip**, the installation program cannot download debugging information from failed installations.

See "Enabling access to the environment" for more information.

- On a command line, create security groups by running the **security-groups.yaml** playbook:

```
$ ansible-playbook -i inventory.yaml security-groups.yaml
```

- On a command line, update the network resources by running the **update-network-resources.yaml** playbook:

```
$ ansible-playbook -i inventory.yaml update-network-resources.yaml 1
```

- 1** This playbook will add tags to the network, subnets, ports, and router. It also attaches floating IP addresses to the API and Ingress ports and sets the security groups for those ports.


5. Optional: If you want to control the default resolvers that Nova servers use, run the RHOSP CLI command:

```
$ openstack subnet set --dns-nameserver <server_1> --dns-nameserver <server_2>
"$INFRA_ID-nodes"
```

6. Optional: You can use the **inventory.yaml** file that you created to customize your installation. For example, you can deploy a cluster that uses bare metal machines.

21.4.17.1. Deploying a cluster with bare metal machines

If you want your cluster to use bare metal machines, modify the **inventory.yaml** file. Your cluster can have both control plane and compute machines running on bare metal, or just compute machines.



NOTE

Be sure that your **install-config.yaml** file reflects whether the RHOSP network that you use for bare metal workers supports floating IP addresses or not.

Prerequisites

- The RHOSP [Bare Metal service \(Ironic\)](#) is enabled and accessible via the RHOSP Compute API.
- Bare metal is available as [a RHOSP flavor](#).
- If your cluster runs on an RHOSP version that is more than 16.1.6 and less than 16.2.4, bare metal workers do not function due to a [known issue](#) that causes the metadata service to be unavailable for services on OpenShift Container Platform nodes.
- The RHOSP network supports both VM and bare metal server attachment.
- Your network configuration does not rely on a provider network. Provider networks are not supported.
- If you want to deploy the machines on a pre-existing network, a RHOSP subnet is provisioned.
- If you want to deploy the machines on an installer-provisioned network, the RHOSP Bare Metal service (Ironic) is able to listen for and interact with Preboot eXecution Environment (PXE) boot machines that run on tenant networks.
- You created an **inventory.yaml** file as part of the OpenShift Container Platform installation process.

Procedure

1. In the **inventory.yaml** file, edit the flavors for machines:
 - a. If you want to use bare-metal control plane machines, change the value of **os_flavor_master** to a bare metal flavor.
 - b. Change the value of **os_flavor_worker** to a bare metal flavor.

An example bare metal inventory.yaml file

```
all:
```

```

hosts:
  localhost:
    ansible_connection: local
    ansible_python_interpreter: "{{ansible_playbook_python}}"

# User-provided values
os_subnet_range: '10.0.0.0/16'
os_flavor_master: 'my-bare-metal-flavor' 1
os_flavor_worker: 'my-bare-metal-flavor' 2
os_image_rhcos: 'rhcos'
os_external_network: 'external'
...

```

- 1 If you want to have bare-metal control plane machines, change this value to a bare metal flavor.
- 2 Change this value to a bare metal flavor to use for compute machines.

Use the updated **inventory.yaml** file to complete the installation process. Machines that are created during deployment use the flavor that you added to the file.



NOTE

The installer may time out while waiting for bare metal machines to boot.

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

21.4.18. Creating the bootstrap machine on RHOSP

Create a bootstrap machine and give it the network access it needs to run on Red Hat OpenStack Platform (RHOSP). Red Hat provides an Ansible playbook that you run to simplify this process.

Prerequisites

- You downloaded the modules in "Downloading playbook dependencies".
- You downloaded the playbooks in "Downloading the installation playbooks".
- The **inventory.yaml**, **common.yaml**, and **bootstrap.yaml** Ansible playbooks are in a common directory.
- The **metadata.json** file that the installation program created is in the same directory as the Ansible playbooks.

Procedure

1. On a command line, change the working directory to the location of the playbooks.
2. On a command line, run the **bootstrap.yaml** playbook:

```
$ ansible-playbook -i inventory.yaml bootstrap.yaml
```

3. After the bootstrap server is active, view the logs to verify that the Ignition files were received:

```
$ openstack console log show "$INFRA_ID-bootstrap"
```

21.4.19. Creating the control plane machines on RHOSP

Create three control plane machines by using the Ignition config files that you generated. Red Hat provides an Ansible playbook that you run to simplify this process.

Prerequisites

- You downloaded the modules in "Downloading playbook dependencies".
- You downloaded the playbooks in "Downloading the installation playbooks".
- The infrastructure ID from the installation program's metadata file is set as an environment variable (**\$INFRA_ID**).
- The **inventory.yaml**, **common.yaml**, and **control-plane.yaml** Ansible playbooks are in a common directory.
- You have the three Ignition files that were created in "Creating control plane Ignition config files".

Procedure

1. On a command line, change the working directory to the location of the playbooks.
2. If the control plane Ignition config files are not already in your working directory, copy them into it.
3. On a command line, run the **control-plane.yaml** playbook:

```
$ ansible-playbook -i inventory.yaml control-plane.yaml
```

4. Run the following command to monitor the bootstrapping process:

```
$ openshift-install wait-for bootstrap-complete
```

You will see messages that confirm that the control plane machines are running and have joined the cluster:

```
INFO API v1.29.4 up
INFO Waiting up to 30m0s for bootstrapping to complete...
...
INFO It is now safe to remove the bootstrap resources
```

21.4.20. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the

correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

21.4.21. Deleting bootstrap resources from RHOSP

Delete the bootstrap resources that you no longer need.

Prerequisites

- You downloaded the modules in "Downloading playbook dependencies".
- You downloaded the playbooks in "Downloading the installation playbooks".
- The **inventory.yaml**, **common.yaml**, and **down-bootstrap.yaml** Ansible playbooks are in a common directory.
- The control plane machines are running.
 - If you do not know the status of the machines, see "Verifying cluster status".

Procedure

1. On a command line, change the working directory to the location of the playbooks.
2. On a command line, run the **down-bootstrap.yaml** playbook:

```
$ ansible-playbook -i inventory.yaml down-bootstrap.yaml
```

The bootstrap port, server, and floating IP address are deleted.

**WARNING**

If you did not disable the bootstrap Ignition file URL earlier, do so now.

21.4.22. Creating compute machines on RHOSP

After standing up the control plane, create compute machines. Red Hat provides an Ansible playbook that you run to simplify this process.

Prerequisites

- You downloaded the modules in "Downloading playbook dependencies".
- You downloaded the playbooks in "Downloading the installation playbooks".
- The **inventory.yaml**, **common.yaml**, and **compute-nodes.yaml** Ansible playbooks are in a common directory.
- The **metadata.json** file that the installation program created is in the same directory as the Ansible playbooks.
- The control plane is active.

Procedure

1. On a command line, change the working directory to the location of the playbooks.
2. On a command line, run the playbook:

```
$ ansible-playbook -i inventory.yaml compute-nodes.yaml
```

Next steps

- Approve the certificate signing requests for the machines.

21.4.23. Approving the certificate signing requests for your machines

When you add machines to a cluster, two pending certificate signing requests (CSRs) are generated for each machine that you added. You must confirm that these CSRs are approved or, if necessary, approve them yourself. The client requests must be approved first, followed by the server requests.

Prerequisites

- You added machines to your cluster.

Procedure

1. Confirm that the cluster recognizes the machines:


```
$ oc get nodes
```

-

Example output

NAME	STATUS	ROLES	AGE	VERSION
master-0	Ready	master	63m	v1.29.4
master-1	Ready	master	63m	v1.29.4
master-2	Ready	master	64m	v1.29.4

The output lists all of the machines that you created.



NOTE

The preceding output might not include the compute nodes, also known as worker nodes, until some CSRs are approved.

- Review the pending CSRs and ensure that you see the client requests with the **Pending** or **Approved** status for each machine that you added to the cluster:


```
$ oc get csr
```

Example output

NAME	AGE	REQUESTOR	CONDITION
csr-8b2br	15m	system:serviceaccount:openshift-machine-config-operator:node-bootstrapper	Pending
csr-8vnps	15m	system:serviceaccount:openshift-machine-config-operator:node-bootstrapper	Pending
...			

In this example, two machines are joining the cluster. You might see more approved CSRs in the list.

- If the CSRs were not approved, after all of the pending CSRs for the machines you added are in **Pending** status, approve the CSRs for your cluster machines:



NOTE

Because the CSRs rotate automatically, approve your CSRs within an hour of adding the machines to the cluster. If you do not approve them within an hour, the certificates will rotate, and more than two certificates will be present for each node. You must approve all of these certificates. After the client CSR is approved, the Kubelet creates a secondary CSR for the serving certificate, which requires manual approval. Then, subsequent serving certificate renewal requests are automatically approved by the **machine-approver** if the Kubelet requests a new certificate with identical parameters.

**NOTE**

For clusters running on platforms that are not machine API enabled, such as bare metal and other user-provisioned infrastructure, you must implement a method of automatically approving the kubelet serving certificate requests (CSRs). If a request is not approved, then the **oc exec**, **oc rsh**, and **oc logs** commands cannot succeed, because a serving certificate is required when the API server connects to the kubelet. Any operation that contacts the Kubelet endpoint requires this certificate approval to be in place. The method must watch for new CSRs, confirm that the CSR was submitted by the **node-bootstrap** service account in the **system:node** or **system:admin** groups, and confirm the identity of the node.

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{\n"}\n{{end}}\n{{end}}' | xargs --no-run-if-empty oc adm certificate approve
```

**NOTE**

Some Operators might not become available until some CSRs are approved.

4. Now that your client requests are approved, you must review the server requests for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

```
NAME      AGE   REQUESTOR                                     CONDITION
csr-bfd72 5m26s system:node:ip-10-0-50-126.us-east-2.compute.internal
Pending
csr-c57lv 5m26s system:node:ip-10-0-95-157.us-east-2.compute.internal
Pending
...
```

5. If the remaining CSRs are not approved, and are in the **Pending** status, approve the CSRs for your cluster machines:

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{"\n"}\n{{end}}' | xargs oc adm certificate approve
```

6. After all client and server CSRs have been approved, the machines have the **Ready** status. Verify this by running the following command:

```
$ oc get nodes
```

Example output

```
NAME      STATUS  ROLES  AGE  VERSION
master-0  Ready   master 73m  v1.29.4
master-1  Ready   master 73m  v1.29.4
master-2  Ready   master 74m  v1.29.4
worker-0  Ready   worker 11m  v1.29.4
worker-1  Ready   worker 11m  v1.29.4
```



NOTE

It can take a few minutes after approval of the server CSRs for the machines to transition to the **Ready** status.

Additional information

- For more information on CSRs, see [Certificate Signing Requests](#).

21.4.24. Verifying a successful installation

Verify that the OpenShift Container Platform installation is complete.

Prerequisites

- You have the installation program (**openshift-install**)

Procedure

- On a command line, enter:

```
$ openshift-install --log-level debug wait-for install-complete
```

The program outputs the console URL, as well as the administrator's login information.

21.4.25. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service

21.4.26. Next steps


- [Customize your cluster.](#)
- If necessary, you can [opt out of remote health reporting](#) .
- If you need to enable external access to node ports, [configure ingress cluster traffic by using a node port.](#)
- If you did not configure RHOSP to accept application traffic over floating IP addresses, [configure RHOSP access with floating IP addresses](#) .

21.5. INSTALLING A CLUSTER ON OPENSTACK IN A RESTRICTED NETWORK

In OpenShift Container Platform 4.16, you can install a cluster on Red Hat OpenStack Platform (RHOSP) in a restricted network by creating an internal mirror of the installation release content.

21.5.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users.](#)
- You verified that OpenShift Container Platform 4.16 is compatible with your RHOSP version by using the [Supported platforms for OpenShift clusters](#) section. You can also compare platform support across different versions by viewing the [OpenShift Container Platform on RHOSP support matrix.](#)
- You [created a registry on your mirror host](#) and obtained the **imageContentSources** data for your version of OpenShift Container Platform.



IMPORTANT

Because the installation media is on the mirror host, you can use that computer to complete all installation steps.

- You understand performance and scalability practices for cluster scaling, control plane sizing, and etcd. For more information, see [Recommended practices for scaling the cluster](#) .
- You have the metadata service enabled in RHOSP.

21.5.2. About installations in restricted networks

In OpenShift Container Platform 4.16, you can perform an installation that does not require an active connection to the internet to obtain software components. Restricted network installations can be completed using installer-provisioned infrastructure or user-provisioned infrastructure, depending on the cloud platform to which you are installing the cluster.

If you choose to perform a restricted network installation on a cloud platform, you still require access to its cloud APIs. Some cloud functions, like Amazon Web Service’s Route 53 DNS and IAM services, require internet access. Depending on your network, you might require less internet access for an installation on bare metal hardware, Nutanix, or on VMware vSphere.

To complete a restricted network installation, you must create a registry that mirrors the contents of the OpenShift image registry and contains the installation media. You can create this registry on a mirror host, which can access both the internet and your closed network, or by using other methods that meet your restrictions.

21.5.2.1. Additional limits

Clusters in restricted networks have the following additional limitations and restrictions:

- The **ClusterVersion** status includes an **Unable to retrieve available updates** error.
- By default, you cannot use the contents of the Developer Catalog because you cannot access the required image stream tags.

21.5.3. Resource guidelines for installing OpenShift Container Platform on RHOSP

To support an OpenShift Container Platform installation, your Red Hat OpenStack Platform (RHOSP) quota must meet the following requirements:

Table 21.5. Recommended resources for a default OpenShift Container Platform cluster on RHOSP

Resource	Value
Floating IP addresses	3
Ports	15
Routers	1
Subnets	1
RAM	88 GB
vCPUs	22
Volume storage	275 GB
Instances	7
Security groups	3
Security group rules	60

Resource	Value
Server groups	2 - plus 1 for each additional availability zone in each machine pool

A cluster might function with fewer than recommended resources, but its performance is not guaranteed.



IMPORTANT

If RHOSP object storage (Swift) is available and operated by a user account with the **swiftoperator** role, it is used as the default backend for the OpenShift Container Platform image registry. In this case, the volume storage requirement is 175 GB. Swift space requirements vary depending on the size of the image registry.



NOTE

By default, your security group and security group rule quotas might be low. If you encounter problems, run **openstack quota set --secgroups 3 --secgroup-rules 60 <project>** as an administrator to increase them.

An OpenShift Container Platform deployment comprises control plane machines, compute machines, and a bootstrap machine.

21.5.3.1. Control plane machines

By default, the OpenShift Container Platform installation process creates three control plane machines.

Each machine requires:

- An instance from the RHOSP quota
- A port from the RHOSP quota
- A flavor with at least 16 GB memory and 4 vCPUs
- At least 100 GB storage space from the RHOSP quota

21.5.3.2. Compute machines

By default, the OpenShift Container Platform installation process creates three compute machines.

Each machine requires:

- An instance from the RHOSP quota
- A port from the RHOSP quota
- A flavor with at least 8 GB memory and 2 vCPUs
- At least 100 GB storage space from the RHOSP quota

TIP

Compute machines host the applications that you run on OpenShift Container Platform; aim to run as many as you can.

21.5.3.3. Bootstrap machine

During installation, a bootstrap machine is temporarily provisioned to stand up the control plane. After the production control plane is ready, the bootstrap machine is deprovisioned.

The bootstrap machine requires:

- An instance from the RHOSP quota
- A port from the RHOSP quota
- A flavor with at least 16 GB memory and 4 vCPUs
- At least 100 GB storage space from the RHOSP quota

21.5.4. Internet access for OpenShift Container Platform


In OpenShift Container Platform 4.16, you require access to the internet to obtain the images that are necessary to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.

21.5.5. Enabling Swift on RHOSP

Swift is operated by a user account with the **swiftoperator** role. Add the role to an account before you run the installation program.



IMPORTANT

If [the Red Hat OpenStack Platform \(RHOSP\) object storage service](#), commonly known as Swift, is available, OpenShift Container Platform uses it as the image registry storage. If it is unavailable, the installation program relies on the RHOSP block storage service, commonly known as Cinder.

If Swift is present and you want to use it, you must enable access to it. If it is not present, or if you do not want to use it, skip this section.



IMPORTANT

RHOSP 17 sets the **rgw_max_attr_size** parameter of Ceph RGW to 256 characters. This setting causes issues with uploading container images to the OpenShift Container Platform registry. You must set the value of **rgw_max_attr_size** to at least 1024 characters.

Before installation, check if your RHOSP deployment is affected by this problem. If it is, reconfigure Ceph RGW.

Prerequisites

- You have a RHOSP administrator account on the target environment.
- The Swift service is installed.
- On [Ceph RGW](#), the **account in url** option is enabled.

Procedure

To enable Swift on RHOSP:

1. As an administrator in the RHOSP CLI, add the **swiftoperator** role to the account that will access Swift:

```
$ openstack role add --user <user> --project <project> swiftoperator
```

Your RHOSP deployment can now use Swift for the image registry.

21.5.6. Defining parameters for the installation program

The OpenShift Container Platform installation program relies on a file that is called **clouds.yaml**. The file describes Red Hat OpenStack Platform (RHOSP) configuration parameters, including the project name, log in information, and authorization service URLs.

Procedure

1. Create the **clouds.yaml** file:
 - If your RHOSP distribution includes the Horizon web UI, generate a **clouds.yaml** file in it.



IMPORTANT

Remember to add a password to the **auth** field. You can also keep secrets in [a separate file](#) from **clouds.yaml**.

- If your RHOSP distribution does not include the Horizon web UI, or you do not want to use Horizon, create the file yourself. For detailed information about **clouds.yaml**, see [Config files](#) in the RHOSP documentation.

```
clouds:
  shiftstack:
    auth:
      auth_url: http://10.10.14.42:5000/v3
      project_name: shiftstack
```

```

username: <username>
password: <password>
user_domain_name: Default
project_domain_name: Default
dev-env:
region_name: RegionOne
auth:
username: <username>
password: <password>
project_name: 'devonly'
auth_url: 'https://10.10.14.22:5001/v2.0'

```

2. If your RHOSP installation uses self-signed certificate authority (CA) certificates for endpoint authentication:
 - a. Copy the certificate authority file to your machine.
 - b. Add the **cacerts** key to the **clouds.yaml** file. The value must be an absolute, non-root-accessible path to the CA certificate:

```

clouds:
  shiftstack:
    ...
  cacert: "/etc/pki/ca-trust/source/anchors/ca.crt.pem"

```

TIP

After you run the installer with a custom CA certificate, you can update the certificate by editing the value of the **ca-cert.pem** key in the **cloud-provider-config** keymap. On a command line, run:

```
$ oc edit configmap -n openshift-config cloud-provider-config
```

3. Place the **clouds.yaml** file in one of the following locations:
 - a. The value of the **OS_CLIENT_CONFIG_FILE** environment variable
 - b. The current directory
 - c. A Unix-specific user configuration directory, for example **~/.config/openstack/clouds.yaml**
 - d. A Unix-specific site configuration directory, for example **/etc/openstack/clouds.yaml**

The installation program searches for **clouds.yaml** in that order.

21.5.7. Setting OpenStack Cloud Controller Manager options

Optionally, you can edit the OpenStack Cloud Controller Manager (CCM) configuration for your cluster. This configuration controls how OpenShift Container Platform interacts with Red Hat OpenStack Platform (RHOSP).

For a complete list of configuration parameters, see the "OpenStack Cloud Controller Manager reference guide" page in the "Installing on OpenStack" documentation.

Procedure

1. If you have not already generated manifest files for your cluster, generate them by running the following command:

```
$ openshift-install --dir <destination_directory> create manifests
```


2. In a text editor, open the cloud-provider configuration manifest file. For example:

```
$ vi openshift/manifests/cloud-provider-config.yaml
```

3. Modify the options according to the CCM reference guide. Configuring Octavia for load balancing is a common case. For example:


```
#...
[LoadBalancer]
lb-provider = "amphora" 1
floating-network-id="d3deb660-4190-40a3-91f1-3732fe6ec4a" 2
create-monitor = True 3
monitor-delay = 10s 4
monitor-timeout = 10s 5
monitor-max-retries = 1 6
#...
```

- 1** This property sets the Octavia provider that your load balancer uses. It accepts **"ovn"** or **"amphora"** as values. If you choose to use OVN, you must also set **lb-method** to **SOURCE_IP_PORT**.
- 2** This property is required if you want to use multiple external networks with your cluster. The cloud provider creates floating IP addresses on the network that is specified here.
- 3** This property controls whether the cloud provider creates health monitors for Octavia load balancers. Set the value to **True** to create health monitors. As of RHOSP 16.2, this feature is only available for the Amphora provider.
- 4** This property sets the frequency with which endpoints are monitored. The value must be in the **time.ParseDuration()** format. This property is required if the value of the **create-monitor** property is **True**.
- 5** This property sets the time that monitoring requests are open before timing out. The value must be in the **time.ParseDuration()** format. This property is required if the value of the **create-monitor** property is **True**.
- 6** This property defines how many successful monitoring requests are required before a load balancer is marked as online. The value must be an integer. This property is required if the value of the **create-monitor** property is **True**.



IMPORTANT

Prior to saving your changes, verify that the file is structured correctly. Clusters might fail if properties are not placed in the appropriate section.



IMPORTANT

You must set the value of the **create-monitor** property to **True** if you use services that have the value of the **.spec.externalTrafficPolicy** property set to **Local**. The OVN Octavia provider in RHOSP 16.2 does not support health monitors. Therefore, services that have **ETP** parameter values set to **Local** might not respond when the **lb-provider** value is set to **"ovn"**.

4. Save the changes to the file and proceed with installation.

TIP

You can update your cloud provider configuration after you run the installer. On a command line, run:

```
$ oc edit configmap -n openshift-config cloud-provider-config
```

After you save your changes, your cluster will take some time to reconfigure itself. The process is complete if none of your nodes have a **SchedulingDisabled** status.

21.5.8. Creating the RHCOS image for restricted network installations

Download the Red Hat Enterprise Linux CoreOS (RHCOS) image to install OpenShift Container Platform on a restricted network Red Hat OpenStack Platform (RHOSP) environment.

Prerequisites

- Obtain the OpenShift Container Platform installation program. For a restricted network installation, the program is on your mirror registry host.

Procedure


1. Log in to the Red Hat Customer Portal's [Product Downloads page](#).
2. Under **Version**, select the most recent release of OpenShift Container Platform 4.16 for RHEL 8.



IMPORTANT

The RHCOS images might not change with every release of OpenShift Container Platform. You must download images with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Use the image versions that match your OpenShift Container Platform version if they are available.

3. Download the **Red Hat Enterprise Linux CoreOS (RHCOS) - OpenStack Image (QCOW)** image.
4. Decompress the image.

**NOTE**

You must decompress the image before the cluster can use it. The name of the downloaded file might not contain a compression extension, like **.gz** or **.tgz**. To find out if or how the file is compressed, in a command line, enter:

```
$ file <name_of_downloaded_file>
```

5. Upload the image that you decompressed to a location that is accessible from the bastion server, like Glance. For example:

```
$ openstack image create --file rhcos-44.81.202003110027-0-openstack.x86_64.qcow2 --
disk-format qcow2 rhcos- $\{RHCOS\_VERSION\}$ 
```

**IMPORTANT**

Depending on your RHOSP environment, you might be able to upload the image in either **.raw** or **.qcow2** formats. If you use Ceph, you must use the **.raw** format.

**WARNING**

If the installation program finds multiple images with the same name, it chooses one of them at random. To avoid this behavior, create unique names for resources in RHOSP.

The image is now available for a restricted installation. Note the image name or location for use in OpenShift Container Platform deployment.

21.5.9. Creating the installation configuration file

You can customize the OpenShift Container Platform cluster you install on Red Hat OpenStack Platform (RHOSP).

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster. For a restricted network installation, these files are on your mirror host.
- You have the **imageContentSources** values that were generated during mirror registry creation.
- You have obtained the contents of the certificate for your mirror registry.
- You have retrieved a Red Hat Enterprise Linux CoreOS (RHCOS) image and uploaded it to an accessible location.

Procedure

1. Create the **install-config.yaml** file.
 - a. Change to the directory that contains the installation program and run the following command:

```
┌ $ ./openshift-install create install-config --dir <installation_directory> 1
```

- 1 For **<installation_directory>**, specify the directory name to store the files that the installation program creates.

When specifying the directory:

- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
 - Use an empty directory. Some installation assets, such as bootstrap X.509 certificates, have short expiration intervals, therefore you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.
- b. At the prompts, provide the configuration details for your cloud:
 - i. Optional: Select an SSH key to use to access your cluster machines.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- ii. Select **openstack** as the platform to target.
 - iii. Specify the Red Hat OpenStack Platform (RHOSP) external network name to use for installing the cluster.
 - iv. Specify the floating IP address to use for external access to the OpenShift API.
 - v. Specify a RHOSP flavor with at least 16 GB RAM to use for control plane nodes and 8 GB RAM for compute nodes.
 - vi. Select the base domain to deploy the cluster to. All DNS records will be sub-domains of this base and will also include the cluster name.
 - vii. Enter a name for your cluster. The name must be 14 or fewer characters long.
2. In the **install-config.yaml** file, set the value of **platform.openstack.clusterOSImage** to the image location or name. For example:

```
┌ platform:
  openstack:
    clusterOSImage: http://mirror.example.com/images/rhcos-43.81.201912131630.0-
    openstack.x86_64.qcow2.gz?
    sha256=ffebbd68e8a1f2a245ca19522c16c86f67f9ac8e4e0c1f0a812b068b16f7265d
```

3. Edit the **install-config.yaml** file to give the additional information that is required for an installation in a restricted network.
 - a. Update the **pullSecret** value to contain the authentication information for your registry:

```
pullSecret: '{"auths":{"<mirror_host_name>:5000": {"auth": "<credentials>","email":
"you@example.com"}}}'
```

For **<mirror_host_name>**, specify the registry domain name that you specified in the certificate for your mirror registry, and for **<credentials>**, specify the base64-encoded user name and password for your mirror registry.

- b. Add the **additionalTrustBundle** parameter and value.

```
additionalTrustBundle: |
  -----BEGIN CERTIFICATE-----

  /-----END CERTIFICATE-----
```

The value must be the contents of the certificate file that you used for your mirror registry. The certificate file can be an existing, trusted certificate authority, or the self-signed certificate that you generated for the mirror registry.

- c. Add the image content resources, which resemble the following YAML excerpt:

```
imageContentSources:
- mirrors:
  - <mirror_host_name>:5000/<repo_name>/release
  source: quay.io/openshift-release-dev/ocp-release
- mirrors:
  - <mirror_host_name>:5000/<repo_name>/release
  source: registry.redhat.io/ocp/release
```


For these values, use the **imageContentSources** that you recorded during mirror registry creation.

- d. Optional: Set the publishing strategy to **Internal**:

```
publish: Internal
```

By setting this option, you create an internal Ingress Controller and a private load balancer.

4. Make any other modifications to the **install-config.yaml** file that you require. For more information about the parameters, see "Installation configuration parameters".
5. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

Additional resources

- [Installation configuration parameters for OpenStack](#)

21.5.9.1. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations.

- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.

**NOTE**

The installation program does not support the proxy **readinessEndpoints** field.

**NOTE**

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$. /openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.

**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

21.5.9.2. Sample customized **install-config.yaml** file for restricted OpenStack installations

This sample **install-config.yaml** demonstrates all of the possible Red Hat OpenStack Platform (RHOSP) customization options.

**IMPORTANT**

This sample file is provided for reference only. You must obtain your **install-config.yaml** file by using the installation program.

```
apiVersion: v1
baseDomain: example.com
controlPlane:
  name: master
  platform: {}
  replicas: 3
compute:
- name: worker
  platform:
    openstack:
      type: ml.large
```

```

replicas: 3
metadata:
  name: example
networking:
  clusterNetwork:
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  machineNetwork:
  - cidr: 10.0.0.0/16
  serviceNetwork:
  - 172.30.0.0/16
  networkType: OVNKubernetes
platform:
  openstack:
    region: region1
    cloud: mycloud
    externalNetwork: external
    computeFlavor: m1.xlarge
    apiFloatingIP: 128.0.0.1
  fips: false
  pullSecret: '{"auths": ...}'
  sshKey: ssh-ed25519 AAAA...
  additionalTrustBundle: |

  -----BEGIN CERTIFICATE-----

  /-----END CERTIFICATE-----

imageContentSources:
- mirrors:
  - <mirror_registry>/<repo_name>/release
  source: quay.io/openshift-release-dev/ocp-release
- mirrors:
  - <mirror_registry>/<repo_name>/release
  source: quay.io/openshift-release-dev/ocp-v4.0-art-dev

```

21.5.10. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.

**IMPORTANT**

Do not skip this procedure in production environments, where disaster recovery and debugging is required.

**NOTE**


You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.

**NOTE**

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.


2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.

**NOTE**

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```

**NOTE**

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

21.5.11. Enabling access to the environment

At deployment, all OpenShift Container Platform machines are created in a Red Hat OpenStack Platform (RHOSP)-tenant network. Therefore, they are not accessible directly in most RHOSP deployments.

You can configure OpenShift Container Platform API and application access by using floating IP addresses (FIPs) during installation. You can also complete an installation without configuring FIPs, but the installer will not configure a way to reach the API or applications externally.

21.5.11.1. Enabling access with floating IP addresses

Create floating IP (FIP) addresses for external access to the OpenShift Container Platform API and cluster applications.

Procedure

1. Using the Red Hat OpenStack Platform (RHOSP) CLI, create the API FIP:

```
$ openstack floating ip create --description "API <cluster_name>.<base_domain>"
<external_network>
```

2. Using the Red Hat OpenStack Platform (RHOSP) CLI, create the apps, or Ingress, FIP:

```
$ openstack floating ip create --description "Ingress <cluster_name>.<base_domain>"
<external_network>
```

3. Add records that follow these patterns to your DNS server for the API and Ingress FIPs:

```
api.<cluster_name>.<base_domain>. IN A <API_FIP>
*.apps.<cluster_name>.<base_domain>. IN A <apps_FIP>
```



NOTE

If you do not control the DNS server, you can access the cluster by adding the cluster domain names such as the following to your `/etc/hosts` file:

- `<api_floating_ip> api.<cluster_name>.<base_domain>`
- `<application_floating_ip> grafana-openshift-monitoring.apps.<cluster_name>.<base_domain>`
- `<application_floating_ip> prometheus-k8s-openshift-monitoring.apps.<cluster_name>.<base_domain>`
- `<application_floating_ip> oauth-openshift.apps.<cluster_name>.<base_domain>`
- `<application_floating_ip> console-openshift-console.apps.<cluster_name>.<base_domain>`
- `application_floating_ip integrated-oauth-server-openshift-authentication.apps.<cluster_name>.<base_domain>`

The cluster domain names in the `/etc/hosts` file grant access to the web console and the monitoring interface of your cluster locally. You can also use the **kubectl** or **oc**. You can access the user applications by using the additional entries pointing to the `<application_floating_ip>`. This action makes the API and applications accessible to only you, which is not suitable for production deployment, but does allow installation for development and testing.

4. Add the FIPs to the `install-config.yaml` file as the values of the following parameters:

- `platform.openstack.ingressFloatingIP`
- `platform.openstack.apiFloatingIP`

If you use these values, you must also enter an external network as the value of the `platform.openstack.externalNetwork` parameter in the `install-config.yaml` file.

TIP

You can make OpenShift Container Platform resources available outside of the cluster by assigning a floating IP address and updating your firewall configuration.

21.5.11.2. Completing installation without floating IP addresses

You can install OpenShift Container Platform on Red Hat OpenStack Platform (RHOSP) without providing floating IP addresses.

In the `install-config.yaml` file, do not define the following parameters:

- `platform.openstack.ingressFloatingIP`

- **platform.openstack.apiFloatingIP**

If you cannot provide an external network, you can also leave **platform.openstack.externalNetwork** blank. If you do not provide a value for **platform.openstack.externalNetwork**, a router is not created for you, and, without additional action, the installer will fail to retrieve an image from Glance. You must configure external connectivity on your own.

If you run the installer from a system that cannot reach the cluster API due to a lack of floating IP addresses or name resolution, installation fails. To prevent installation failure in these cases, you can use a proxy network or run the installer from a system that is on the same network as your machines.



NOTE

You can enable name resolution by creating DNS records for the API and Ingress ports. For example:

```
api.<cluster_name>.<base_domain>. IN A <api_port_IP>
*.apps.<cluster_name>.<base_domain>. IN A <ingress_port_IP>
```

If you do not control the DNS server, you can add the record to your **/etc/hosts** file. This action makes the API accessible to only you, which is not suitable for production deployment but does allow installation for development and testing.

21.5.12. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.

Procedure

- Change to the directory that contains the installation program and initialize the cluster deployment:


```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

- 1 For **<installation_directory>**, specify the location of your customized **./install-config.yaml** file.
- 2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to **<installation_directory>/openshift_install.log**.




IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

21.5.13. Verifying cluster status

You can verify your OpenShift Container Platform cluster's status during or after installation.

Procedure

1. In the cluster environment, export the administrator's kubeconfig file:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server.

2. View the control plane and compute machines created after a deployment:

```
$ oc get nodes
```

3. View your cluster's version:

```
$ oc get clusterversion
```

4. View your Operators' status:

```
$ oc get clusteroperator
```

5. View all running pods in the cluster:

```
$ oc get pods -A
```

21.5.14. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

Additional resources

- See [Accessing the web console](#) for more details about accessing and understanding the OpenShift Container Platform web console.

21.5.15. Disabling the default OperatorHub catalog sources

Operator catalogs that source content provided by Red Hat and community projects are configured for OperatorHub by default during an OpenShift Container Platform installation. In a restricted network environment, you must disable the default catalogs as a cluster administrator.

Procedure

- Disable the sources for the default catalogs by adding **disableAllDefaultSources: true** to the **OperatorHub** object:

```
$ oc patch OperatorHub cluster --type json \
  -p '[{"op": "add", "path": "/spec/disableAllDefaultSources", "value": true}]'
```

TIP

Alternatively, you can use the web console to manage catalog sources. From the **Administration** → **Cluster Settings** → **Configuration** → **OperatorHub** page, click the **Sources** tab, where you can create, update, delete, disable, and enable individual sources.

21.5.16. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service

21.5.17. Next steps

- [Customize your cluster](#).
- If the mirror registry that you used to install your cluster has a trusted CA, add it to the cluster by [configuring additional trust stores](#).
- If necessary, you can [opt out of remote health reporting](#).
- If necessary, see [Registering your disconnected cluster](#)
- [Configure image streams](#) for the Cluster Samples Operator and the **must-gather** tool.
- Learn how to [use Operator Lifecycle Manager \(OLM\) on restricted networks](#).
- If you did not configure RHOSP to accept application traffic over floating IP addresses, [configure RHOSP access with floating IP addresses](#).

21.6. OPENSTACK CLOUD CONTROLLER MANAGER REFERENCE GUIDE

21.6.1. The OpenStack Cloud Controller Manager

Beginning with OpenShift Container Platform 4.12, clusters that run on Red Hat OpenStack Platform (RHOSP) were switched from the legacy OpenStack cloud provider to the external OpenStack Cloud Controller Manager (CCM). This change follows the move in Kubernetes from in-tree, legacy cloud providers to external cloud providers that are implemented by using the [Cloud Controller Manager](#).

To preserve user-defined configurations for the legacy cloud provider, existing configurations are mapped to new ones as part of the migration process. It searches for a configuration called **cloud-provider-config** in the **openshift-config** namespace.



NOTE

The config map name **cloud-provider-config** is not statically configured. It is derived from the **spec.cloudConfig.name** value in the **infrastructure/cluster** CRD.

Found configurations are synchronized to the **cloud-conf** config map in the **openshift-cloud-controller-manager** namespace.

As part of this synchronization, the OpenStack CCM Operator alters the new config map such that its properties are compatible with the external cloud provider. The file is changed in the following ways:

- The **[Global] secret-name**, **[Global] secret-namespace**, and **[Global] kubeconfig-path** options are removed. They do not apply to the external cloud provider.
- The **[Global] use-clouds**, **[Global] clouds-file**, and **[Global] cloud** options are added.
- The entire **[BlockStorage]** section is removed. External cloud providers no longer perform storage operations. Block storage configuration is managed by the Cinder CSI driver.

Additionally, the CCM Operator enforces a number of default options. Values for these options are always overridden as follows:


```
[Global]
use-clouds = true
clouds-file = /etc/openstack/secret/clouds.yaml
cloud = openstack
...

[LoadBalancer]
enabled = true
```

The **clouds-value** value, **/etc/openstack/secret/clouds.yaml**, is mapped to the **openstack-cloud-credentials** config in the **openshift-cloud-controller-manager** namespace. You can modify the RHOSP cloud in this file as you do any other **clouds.yaml** file.

21.6.2. The OpenStack Cloud Controller Manager (CCM) config map

An OpenStack CCM config map defines how your cluster interacts with your RHOSP cloud. By default, this configuration is stored under the **cloud.conf** key in the **cloud-conf** config map in the **openshift-cloud-controller-manager** namespace.



IMPORTANT

The **cloud-conf** config map is generated from the **cloud-provider-config** config map in the **openshift-config** namespace.

To change the settings that are described by the **cloud-conf** config map, modify the **cloud-provider-config** config map.

As part of this synchronization, the CCM Operator overrides some options. For more information, see "The RHOSP Cloud Controller Manager".

For example:

An example cloud-conf config map


```
apiVersion: v1
data:
  cloud.conf: |
    [Global] 1
    secret-name = openstack-credentials
    secret-namespace = kube-system
    region = regionOne
    [LoadBalancer]
    enabled = True
kind: ConfigMap
metadata:
  creationTimestamp: "2022-12-20T17:01:08Z"
  name: cloud-conf
  namespace: openshift-cloud-controller-manager
  resourceVersion: "2519"
  uid: cbbeedaf-41ed-41c2-9f37-4885732d3677
```

1 Set global options by using a **clouds.yaml** file rather than modifying the config map.

The following options are present in the config map. Except when indicated otherwise, they are mandatory for clusters that run on RHOSP.

21.6.2.1. Load balancer options

CCM supports several load balancer options for deployments that use Octavia.



NOTE

Neutron-LBaaS support is deprecated.

Option	Description
enabled	Whether or not to enable the LoadBalancer type of services integration. The default value is true .

Option	Description
floating-network-id	<p>Optional. The external network used to create floating IP addresses for load balancer virtual IP addresses (VIPs). If there are multiple external networks in the cloud, this option must be set or the user must specify loadbalancer.openstack.org/floating-network-id in the service annotation.</p>
floating-subnet-id	<p>Optional. The external network subnet used to create floating IP addresses for the load balancer VIP. Can be overridden by the service annotation loadbalancer.openstack.org/floating-subnet-id.</p>
floating-subnet	<p>Optional. A name pattern (glob or regular expression if starting with ~) for the external network subnet used to create floating IP addresses for the load balancer VIP. Can be overridden by the service annotation loadbalancer.openstack.org/floating-subnet. If multiple subnets match the pattern, the first one with available IP addresses is used.</p>
floating-subnet-tags	<p>Optional. Tags for the external network subnet used to create floating IP addresses for the load balancer VIP. Can be overridden by the service annotation loadbalancer.openstack.org/floating-subnet-tags. If multiple subnets match these tags, the first one with available IP addresses is used.</p> <p>If the RHOSP network is configured with sharing disabled, for example, with the --no-share flag used during creation, this option is unsupported. Set the network to share to use this option.</p>

Option	Description
lb-method	<p>The load balancing algorithm used to create the load balancer pool. For the Amphora provider the value can be ROUND_ROBIN, LEAST_CONNECTIONS, or SOURCE_IP. The default value is ROUND_ROBIN.</p> <p>For the OVN provider, only the SOURCE_IP_PORT algorithm is supported.</p> <p>For the Amphora provider, if using the LEAST_CONNECTIONS or SOURCE_IP methods, configure the create-monitor option as true in the cloud-provider-config config map on the openshift-config namespace and ETP:Local on the load-balancer type service to allow balancing algorithm enforcement in the client to service endpoint connections.</p>
lb-provider	<p>Optional. Used to specify the provider of the load balancer, for example, amphora or octavia. Only the Amphora and Octavia providers are supported.</p>
lb-version	<p>Optional. The load balancer API version. Only "v2" is supported.</p>
subnet-id	<p>The ID of the Networking service subnet on which load balancer VIPs are created. For dual stack deployments, leave this option unset. The OpenStack cloud provider automatically selects which subnet to use for a load balancer.</p>
network-id	<p>The ID of the Networking service network on which load balancer VIPs are created. Unnecessary if subnet-id is set. If this property is not set, the network is automatically selected based on the network that cluster nodes use.</p>
create-monitor	<p>Whether or not to create a health monitor for the service load balancer. A health monitor is required for services that declare externalTrafficPolicy: Local. The default value is false.</p> <p>This option is unsupported if you use RHOSP earlier than version 17 with the ovn provider.</p>
monitor-delay	<p>The interval in seconds by which probes are sent to members of the load balancer. The default value is 5.</p>

Option	Description
monitor-max-retries	The number of successful checks that are required to change the operating status of a load balancer member to ONLINE . The valid range is 1 to 10 , and the default value is 1 .
monitor-timeout	The time in seconds that a monitor waits to connect to the back end before it times out. The default value is 3 .
internal-lb	Whether or not to create an internal load balancer without floating IP addresses. The default value is false .
LoadBalancerClass "ClassName"	<p>This is a config section that comprises a set of options:</p> <ul style="list-style-type: none"> ● floating-network-id ● floating-subnet-id ● floating-subnet ● floating-subnet-tags ● network-id ● subnet-id <p>The behavior of these options is the same as that of the identically named options in the load balancer section of the CCM config file.</p> <p>You can set the ClassName value by specifying the service annotation loadbalancer.openstack.org/class.</p>
max-shared-lb	The maximum number of services that can share a load balancer. The default value is 2 .

21.6.2.2. Options that the Operator overrides


The CCM Operator overrides the following options, which you might recognize from configuring RHOSP. Do not configure them yourself. They are included in this document for informational purposes only.

Option	Description
--------	-------------

Option	Description
auth-url	The RHOSP Identity service URL. For example, http://128.110.154.166/identity .
os-endpoint-type	The type of endpoint to use from the service catalog.
username	The Identity service user name.
password	The Identity service user password.
domain-id	The Identity service user domain ID.
domain-name	The Identity service user domain name.
tenant-id	<p>The Identity service project ID. Leave this option unset if you are using Identity service application credentials.</p> <p>In version 3 of the Identity API, which changed the identifier tenant to project, the value of tenant-id is automatically mapped to the project construct in the API.</p>
tenant-name	The Identity service project name.
tenant-domain-id	The Identity service project domain ID.
tenant-domain-name	The Identity service project domain name.
user-domain-id	The Identity service user domain ID.
user-domain-name	The Identity service user domain name.
use-clouds	<p>Whether or not to fetch authorization credentials from a clouds.yaml file. Options set in this section are prioritized over values read from the clouds.yaml file.</p> <p>CCM searches for the file in the following places:</p> <ol style="list-style-type: none"> 1. The value of the clouds-file option. 2. A file path stored in the environment variable OS_CLIENT_CONFIG_FILE. 3. The directory pkg/openstack. 4. The directory ~/config/openstack. 5. The directory /etc/openstack.

Option	Description
clouds-file	The file path of a clouds.yaml file. It is used if the use-clouds option is set to true .
cloud	The named cloud in the clouds.yaml file that you want to use. It is used if the use-clouds option is set to true .

21.7. DEPLOYING ON OPENSTACK WITH ROOTVOLUME AND ETCD ON LOCAL DISK



IMPORTANT

Deploying on Red Hat OpenStack Platform (RHOSP) with rootVolume and etcd on local disk is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

As a day 2 operation, you can resolve and prevent performance issues of your Red Hat OpenStack Platform (RHOSP) installation by moving etcd from a root volume (provided by OpenStack Cinder) to a dedicated ephemeral local disk.

21.7.1. Deploying RHOSP on local disk

If you have an existing RHOSP cloud, you can move etcd from that cloud to a dedicated ephemeral local disk.



WARNING

This procedure is for testing etcd on a local disk only and should not be used on production clusters. In certain cases, complete loss of the control plane can occur. For more information, see "Overview of backup and restore operation" under "Backup and restore".

Prerequisites

- You have an OpenStack cloud with a working Cinder.

- Your OpenStack cloud has at least 75 GB of available storage to accommodate 3 root volumes for the OpenShift control plane.
- The OpenStack cloud is deployed with Nova ephemeral storage that uses a local storage backend and not **rbd**.

Procedure

1. Create a Nova flavor for the control plane with at least 10 GB of ephemeral disk by running the following command, replacing the values for **--ram**, **--disk**, and **<flavor_name>** based on your environment:

```
$ openstack flavor create --<ram 16384> --<disk 0> --ephemeral 10 --vcpus 4
<flavor_name>
```

2. Deploy a cluster with root volumes for the control plane; for example:

Example YAML file

```
# ...
controlPlane:
  name: master
  platform:
    openstack:
      type: ${CONTROL_PLANE_FLAVOR}
      rootVolume:
        size: 25
        types:
          - ${CINDER_TYPE}
      replicas: 3
# ...
```

3. Deploy the cluster you created by running the following command:

```
$ openshift-install create cluster --dir <installation_directory> ❶
```

- ❶ For **<installation_directory>**, specify the location of the customized **./install-config.yaml** file that you previously created.

4. Verify that the cluster you deployed is healthy before proceeding to the next step by running the following command:

```
$ oc wait clusteroperators --all --for=condition=Progressing=false ❶
```

- ❶ Ensures that the cluster operators are finished progressing and that the cluster is not deploying or updating.

5. Edit the **ControlPlaneMachineSet** (CPMS) to add the additional block ephemeral device that is used by etcd by running the following command:

```
$ oc patch ControlPlaneMachineSet/cluster -n openshift-machine-api --type json -p ' ❶
[
```

```

    {
      "op": "add",
      "path":
"/spec/template/machines_v1beta1_machine_openshift_io/spec/providerSpec/value/additionalBlockDevices", ❷
      "value": [
        {
          "name": "etcd",
          "sizeGiB": 10,
          "storage": {
            "type": "Local" ❸
          }
        }
      ]
    }
  ],
],

```

- ❶ Applies the JSON patch to the **ControlPlaneMachineSet** custom resource (CR).
- ❷ Specifies the path where the **additionalBlockDevices** are added.
- ❸ Adds the etcd devices with at least local storage of 10 GB to the cluster. You can specify values greater than 10 GB as long as the etcd device fits the Nova flavor. For example, if the Nova flavor has 15 GB, you can create the etcd device with 12 GB.

6. Verify that the control plane machines are healthy by using the following steps:

- a. Wait for the control plane machine set update to finish by running the following command:

```
$ oc wait --timeout=90m --for=condition=Progressing=false
controlplanemachineset.machine.openshift.io -n openshift-machine-api cluster
```

- b. Verify that the 3 control plane machine sets are updated by running the following command:

```
$ oc wait --timeout=90m --for=jsonpath='{.status.updatedReplicas}'=3
controlplanemachineset.machine.openshift.io -n openshift-machine-api cluster
```

- c. Verify that the 3 control plane machine sets are healthy by running the following command:

```
$ oc wait --timeout=90m --for=jsonpath='{.status.replicas}'=3
controlplanemachineset.machine.openshift.io -n openshift-machine-api cluster
```

- d. Verify that the **ClusterOperators** are not progressing in the cluster by running the following command:

```
$ oc wait clusteroperators --timeout=30m --all --for=condition=Progressing=false
```

- e. Verify that each of the 3 control plane machines has the additional block device you previously created by running the following script:

```
$ cp_machines=$(oc get machines -n openshift-machine-api --
selector='machine.openshift.io/cluster-api-machine-role=master' --no-headers -o custom-
```

```

columns=NAME:.metadata.name) 1

if [[ $(echo "${cp_machines}" | wc -l) -ne 3 ]]; then
  exit 1
fi 2

for machine in ${cp_machines}; do
  if ! oc get machine -n openshift-machine-api "${machine}" -o
jsonpath='{.spec.providerSpec.value.additionalBlockDevices}' | grep -q 'etcd'; then
  exit 1
fi 3
done

```

- 1** Retrieves the control plane machines running in the cluster.
- 2** Iterates over machines which have an **additionalBlockDevices** entry with the name **etcd**.
- 3** Outputs the name of every control plane machine which has an **additionalBlockDevice** named **etcd**.

7. Create a file named **98-var-lib-etcd.yaml** by using the following YAML file:



WARNING

This procedure is for testing etcd on a local disk and should not be used on a production cluster. In certain cases, complete loss of the control plane can occur. For more information, see "Overview of backup and restore operation" under "Backup and restore".

```

apiVersion: machineconfiguration.openshift.io/v1
kind: MachineConfig
metadata:
  labels:
    machineconfiguration.openshift.io/role: master
  name: 98-var-lib-etcd
spec:
  config:
    ignition:
      version: 3.4.0
    systemd:
      units:
        - contents: |
            [Unit]
            Description=Mount local-etcd to /var/lib/etcd

            [Mount]
            What=/dev/disk/by-label/local-etcd 1

```

```

Where=/var/lib/etcd
Type=xfv
Options=defaults,prjquota

[Install]
WantedBy=local-fs.target
enabled: true
name: var-lib-etcd.mount
- contents: |
  [Unit]
  Description=Create local-etcd filesystem
  DefaultDependencies=no
  After=local-fs-pre.target
  ConditionPathIsSymbolicLink=!/dev/disk/by-label/local-etcd 2

  [Service]
  Type=oneshot
  RemainAfterExit=yes
  ExecStart=/bin/bash -c "[ -L /dev/disk/by-label/ephemeral0 ] || ( >&2 echo Ephemeral
disk does not exist; /usr/bin/false )"
  ExecStart=/usr/sbin/mkfs.xfs -f -L local-etcd /dev/disk/by-label/ephemeral0 3

  [Install]
  RequiredBy=dev-disk-by\x2dlabel-local\x2detcd.device
  enabled: true
  name: create-local-etcd.service
- contents: |
  [Unit]
  Description=Migrate existing data to local etcd
  After=var-lib-etcd.mount
  Before=crio.service 4

  Requisite=var-lib-etcd.mount
  ConditionPathExists=!/var/lib/etcd/member
  ConditionPathIsDirectory=/sysroot/ostree/deploy/rhcos/var/lib/etcd/member 5

  [Service]
  Type=oneshot
  RemainAfterExit=yes

  ExecStart=/bin/bash -c "if [ -d /var/lib/etcd/member.migrate ]; then rm -rf
/var/lib/etcd/member.migrate; fi" 6

  ExecStart=/usr/bin/cp -aZ /sysroot/ostree/deploy/rhcos/var/lib/etcd/member/
/var/lib/etcd/member.migrate
  ExecStart=/usr/bin/mv /var/lib/etcd/member.migrate /var/lib/etcd/member 7

  [Install]
  RequiredBy=var-lib-etcd.mount
  enabled: true
  name: migrate-to-local-etcd.service
- contents: |
  [Unit]
  Description=Relabel /var/lib/etcd

```

```
After=migrate-to-local-etcd.service
Before=crio.service
```

```
[Service]
Type=oneshot
RemainAfterExit=yes
```

```
ExecCondition=/bin/bash -c "[ -n \"$(restorecon -nv /var/lib/etcd)\" ]" 8
```

```
ExecStart=/usr/sbin/restorecon -R /var/lib/etcd
```

```
[Install]
RequiredBy=var-lib-etcd.mount
enabled: true
name: relabel-var-lib-etcd.service
```

- 1** The etcd database must be mounted by the device, not a label, to ensure that **systemd** generates the device dependency used in this config to trigger filesystem creation.
- 2** Do not run if the file system **dev/disk/by-label/local-etcd** already exists.
- 3** Fails with an alert message if **/dev/disk/by-label/ephemeral0** does not exist.
- 4** Migrates existing data to local etcd database. This config does so after **/var/lib/etcd** is mounted, but before CRI-O starts so etcd is not running yet.
- 5** Requires that etcd is mounted and does not contain a member directory, but the ostree does.
- 6** Cleans up any previous migration state.
- 7** Copies and moves in separate steps to ensure atomic creation of a complete member directory.
- 8** Performs a quick check of the mount point directory before performing a full recursive relabel. If restorecon in the file path **/var/lib/etcd** cannot rename the directory, the recursive rename is not performed.

8. Create the new **MachineConfig** object by running the following command:

```
$ oc create -f 98-var-lib-etcd.yaml
```



NOTE

Moving the etcd database onto the local disk of each control plane machine takes time.

9. Verify that the etcd databases has been transferred to the local disk of each control plane by running the following commands:
 - a. Verify that the cluster is still updating by running the following command:

```
$ oc wait --timeout=45m --for=condition=Updating=false machineconfigpool/master
```

- b. Verify that the cluster is ready by running the following command:

```
$ oc wait node --selector='node-role.kubernetes.io/master' --for condition=Ready --timeout=30s
```

- c. Verify that the cluster Operators are running in the cluster by running the following command:

```
$ oc wait clusteroperators --timeout=30m --all --for=condition=Progressing=false
```

21.7.2. Additional resources

- [Recommended etcd practices](#)
- [Overview of backup and restore options](#)

21.8. UNINSTALLING A CLUSTER ON OPENSTACK

You can remove a cluster that you deployed to Red Hat OpenStack Platform (RHOSP).

21.8.1. Removing a cluster that uses installer-provisioned infrastructure

You can remove a cluster that uses installer-provisioned infrastructure from your cloud.



NOTE

After uninstallation, check your cloud provider for any resources not removed properly, especially with User Provisioned Infrastructure (UPI) clusters. There might be resources that the installer did not create or that the installer is unable to access.

Prerequisites

- You have a copy of the installation program that you used to deploy the cluster.
- You have the files that the installation program generated when you created your cluster.

Procedure

1. From the directory that contains the installation program on the computer that you used to install the cluster, run the following command:

```
$. /openshift-install destroy cluster \  
--dir <installation_directory> --log-level info 1 2
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.
- 2** To view different details, specify **warn**, **debug**, or **error** instead of **info**.

**NOTE**

You must specify the directory that contains the cluster definition files for your cluster. The installation program requires the **metadata.json** file in this directory to delete the cluster.

2. Optional: Delete the **<installation_directory>** directory and the OpenShift Container Platform installation program.

21.9. UNINSTALLING A CLUSTER ON RHOSP FROM YOUR OWN INFRASTRUCTURE

You can remove a cluster that you deployed to Red Hat OpenStack Platform (RHOSP) on user-provisioned infrastructure.

21.9.1. Downloading playbook dependencies

The Ansible playbooks that simplify the removal process on user-provisioned infrastructure require several Python modules. On the machine where you will run the process, add the modules' repositories and then download them.

**NOTE**

These instructions assume that you are using Red Hat Enterprise Linux (RHEL) 8.

Prerequisites

- Python 3 is installed on your machine.

Procedure

1. On a command line, add the repositories:

- a. Register with Red Hat Subscription Manager:

```
$ sudo subscription-manager register # If not done already
```

- b. Pull the latest subscription data:

```
$ sudo subscription-manager attach --pool=$YOUR_POOLID # If not done already
```

- c. Disable the current repositories:

```
$ sudo subscription-manager repos --disable=* # If not done already
```

- d. Add the required repositories:

```
$ sudo subscription-manager repos \
  --enable=rhel-8-for-x86_64-baseos-rpms \
  --enable=openstack-16-tools-for-rhel-8-x86_64-rpms \
  --enable=ansible-2.9-for-rhel-8-x86_64-rpms \
  --enable=rhel-8-for-x86_64-appstream-rpms
```

2. Install the modules:

```
$ sudo yum install python3-openstackclient ansible python3-openstacksdk
```

3. Ensure that the **python** command points to **python3**:

```
$ sudo alternatives --set python /usr/bin/python3
```

21.9.2. Removing a cluster from RHOSP that uses your own infrastructure

You can remove an OpenShift Container Platform cluster on Red Hat OpenStack Platform (RHOSP) that uses your own infrastructure. To complete the removal process quickly, run several Ansible playbooks.

Prerequisites

- Python 3 is installed on your machine.
- You downloaded the modules in "Downloading playbook dependencies."
- You have the playbooks that you used to install the cluster.
- You modified the playbooks that are prefixed with **down-** to reflect any changes that you made to their corresponding installation playbooks. For example, changes to the **bootstrap.yaml** file are reflected in the **down-bootstrap.yaml** file.
- All of the playbooks are in a common directory.

Procedure

1. On a command line, run the playbooks that you downloaded:

```
$ ansible-playbook -i inventory.yaml \
  down-bootstrap.yaml \
  down-control-plane.yaml \
  down-compute-nodes.yaml \
  down-load-balancers.yaml \
  down-network.yaml \
  down-security-groups.yaml
```

2. Remove any DNS record changes you made for the OpenShift Container Platform installation.


OpenShift Container Platform is removed from your infrastructure.

21.10. INSTALLATION CONFIGURATION PARAMETERS FOR OPENSTACK

Before you deploy an OpenShift Container Platform cluster on Red Hat OpenStack Platform (RHOSP), you provide parameters to customize your cluster and the platform that hosts it. When you create the **install-config.yaml** file, you provide values for the required parameters through the command line. You can then modify the **install-config.yaml** file to customize your cluster further.

21.10.1. Available installation configuration parameters for OpenStack

The following tables specify the required, optional, and OpenStack-specific installation configuration parameters that you can set as part of the installation process.



NOTE

After installation, you cannot modify these parameters in the **install-config.yaml** file.

21.10.1.1. Required configuration parameters

Required installation configuration parameters are described in the following table:

Table 21.6. Required parameters

Parameter	Description	Values
apiVersion:	The API version for the install-config.yaml content. The current version is v1 . The installation program may also support older API versions.	String
baseDomain:	The base domain of your cloud provider. The base domain is used to create routes to your OpenShift Container Platform cluster components. The full DNS name for your cluster is a combination of the baseDomain and metadata.name parameter values that uses the <metadata.name>.<baseDomain> format.	A fully-qualified domain or subdomain name, such as example.com .
metadata:	Kubernetes resource ObjectMeta , from which only the name parameter is consumed.	Object
metadata: name:	The name of the cluster. DNS records for the cluster are all subdomains of {{.metadata.name}}.{{.baseDomain}} .	String of lowercase letters, hyphens (-), and periods (.), such as dev . The string must be 14 characters or fewer long.

Parameter	Description	Values
<code>platform:</code>	The configuration for the specific platform upon which to perform the installation: aws , baremetal , azure , gcp , ibmcloud , nutanix , openstack , powervs , vsphere , or <code>{}</code> . For additional information about platform . <platform> parameters, consult the table for your specific platform that follows.	Object
<code>pullSecret:</code>	Get a pull secret from Red Hat OpenShift Cluster Manager to authenticate downloading container images for OpenShift Container Platform components from services such as Quay.io.	<pre>{ "auths":{ "cloud.openshift.com":{ "auth":"b3Blb=", "email":"you@example.com" }, "quay.io":{ "auth":"b3Blb=", "email":"you@example.com" } } }</pre>

21.10.1.2. Network configuration parameters

You can customize your installation configuration based on the requirements of your existing network infrastructure. For example, you can expand the IP address block for the cluster network or provide different IP address blocks than the defaults.

Only IPv4 addresses are supported.





NOTE

Globalnet is not supported with Red Hat OpenShift Data Foundation disaster recovery solutions. For regional disaster recovery scenarios, ensure that you use a nonoverlapping range of private IP addresses for the cluster and service networks in each cluster.

Table 21.7. Network parameters

Parameter	Description	Values
-----------	-------------	--------

Parameter	Description	Values
<code>networking:</code>	The configuration for the cluster network.	Object  NOTE You cannot modify parameters specified by the networking object after installation.
<code>networking: networkType:</code>	The Red Hat OpenShift Networking network plugin to install.	OVNKubernetes. OVNKubernetes is a CNI plugin for Linux networks and hybrid networks that contain both Linux and Windows servers. The default value is OVNKubernetes .
<code>networking: clusterNetwork:</code>	The IP address blocks for pods. The default value is 10.128.0.0/14 with a host prefix of /23 . If you specify multiple IP address blocks, the blocks must not overlap.	An array of objects. For example: <code>networking: clusterNetwork: - cidr: 10.128.0.0/14 hostPrefix: 23</code>
<code>networking: clusterNetwork: cidr:</code>	Required if you use networking.clusterNetwork . An IP address block. An IPv4 network.	An IP address block in Classless Inter-Domain Routing (CIDR) notation. The prefix length for an IPv4 block is between 0 and 32 .
<code>networking: clusterNetwork: hostPrefix:</code>	The subnet prefix length to assign to each individual node. For example, if hostPrefix is set to 23 then each node is assigned a /23 subnet out of the given cidr . A hostPrefix value of 23 provides 510 ($2^{(32 - 23)} - 2$) pod IP addresses.	A subnet prefix. The default value is 23 .
<code>networking: serviceNetwork:</code>	The IP address block for services. The default value is 172.30.0.0/16 . The OVN-Kubernetes network plugins supports only a single IP address block for the service network.	An array with an IP address block in CIDR format. For example: <code>networking: serviceNetwork: - 172.30.0.0/16</code>

Parameter	Description	Values
<code>networking: machineNetwork:</code>	The IP address blocks for machines. If you specify multiple IP address blocks, the blocks must not overlap.	An array of objects. For example: <code>networking: machineNetwork: - cidr: 10.0.0.0/16</code>
<code>networking: machineNetwork: cidr:</code>	Required if you use networking.machineNetwork . An IP address block. The default value is 10.0.0.0/16 for all platforms other than libvirt and IBM Power® Virtual Server. For libvirt, the default value is 192.168.126.0/24 . For IBM Power® Virtual Server, the default value is 192.168.0.0/24 .	An IP network block in CIDR notation. For example, 10.0.0.0/16 .  NOTE Set the networking.machineNetwork to match the CIDR that the preferred NIC resides in.


21.10.1.3. Optional configuration parameters


Optional installation configuration parameters are described in the following table:

Table 21.8. Optional parameters

Parameter	Description	Values
<code>additionalTrustBundle:</code>	A PEM-encoded X.509 certificate bundle that is added to the nodes' trusted certificate store. This trust bundle may also be used when a proxy has been configured.	String
<code>capabilities:</code>	Controls the installation of optional core cluster components. You can reduce the footprint of your OpenShift Container Platform cluster by disabling optional components. For more information, see the "Cluster capabilities" page in <i>Installing</i> .	String array


Parameter	Description	Values
capabilities: baselineCapabilitySet:	Selects an initial set of optional capabilities to enable. Valid values are None , v4.11 , v4.12 and vCurrent . The default value is vCurrent .	String
capabilities: additionalEnabledCapabilities:	Extends the set of optional capabilities beyond what you specify in baselineCapabilitySet . You may specify multiple capabilities in this parameter.	String array
cpuPartitioningMode:	Enables workload partitioning, which isolates OpenShift Container Platform services, cluster management workloads, and infrastructure pods to run on a reserved set of CPUs. Workload partitioning can only be enabled during installation and cannot be disabled after installation. While this field enables workload partitioning, it does not configure workloads to use specific CPUs. For more information, see the <i>Workload partitioning</i> page in the <i>Scalability and Performance</i> section.	None or AllNodes . None is the default value.
compute:	The configuration for the machines that comprise the compute nodes.	Array of MachinePool objects.
compute: architecture:	Determines the instruction set architecture of the machines in the pool. Currently, clusters with varied architectures are not supported. All pools must specify the same architecture. Valid values are amd64 (the default).	String

Parameter	Description	Values
<code>compute: hyperthreading:</code>	<p>Whether to enable or disable simultaneous multithreading, or hyperthreading, on compute machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.</p>  <p>IMPORTANT</p> <p>If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.</p>	Enabled or Disabled
<code>compute: name:</code>	Required if you use compute . The name of the machine pool.	worker
<code>compute: platform:</code>	Required if you use compute . Use this parameter to specify the cloud provider to host the worker machines. This parameter value must match the controlPlane.platform parameter value.	aws, azure, gcp, ibmcloud, nutanix, openstack, powervs, vsphere, or {}
<code>compute: replicas:</code>	The number of compute machines, which are also known as worker machines, to provision.	A positive integer greater than or equal to 2 . The default value is 3 .
<code>featureSet:</code>	Enables the cluster for a feature set. A feature set is a collection of OpenShift Container Platform features that are not enabled by default. For more information about enabling a feature set during installation, see "Enabling features using feature gates".	String. The name of the feature set to enable, such as TechPreviewNoUpgrade .
<code>controlPlane:</code>	The configuration for the machines that comprise the control plane.	Array of MachinePool objects.

Parameter	Description	Values
<code>controlPlane: architecture:</code>	Determines the instruction set architecture of the machines in the pool. Currently, clusters with varied architectures are not supported. All pools must specify the same architecture. Valid values are amd64 (the default).	String
<code>controlPlane: hyperthreading:</code>	Whether to enable or disable simultaneous multithreading, or hyperthreading , on control plane machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.  IMPORTANT If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.	Enabled or Disabled
<code>controlPlane: name:</code>	Required if you use controlPlane . The name of the machine pool.	master
<code>controlPlane: platform:</code>	Required if you use controlPlane . Use this parameter to specify the cloud provider that hosts the control plane machines. This parameter value must match the compute.platform parameter value.	aws, azure, gcp, ibmcloud, nutanix, openstack, powervs, vsphere , or {}
<code>controlPlane: replicas:</code>	The number of control plane machines to provision.	Supported values are 3 , or 1 when deploying single-node OpenShift.
<code>credentialsMode:</code>	The Cloud Credential Operator (CCO) mode. If no mode is specified, the CCO dynamically tries to determine the capabilities of the provided credentials, with a preference for mint mode on the platforms where multiple modes are supported.	Mint, Passthrough, Manual or an empty string (<code>""</code>). ^[1]

Parameter	Description	Values
<p><code>fips</code></p>	<p>Enable or disable FIPS mode. The default is <code>false</code> (disabled). If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.</p> <div data-bbox="488 477 592 1525" style="background-color: black; color: white; padding: 5px; margin-bottom: 10px;"> <p>IMPORTANT</p> <p>To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see Installing the system in FIPS mode.</p> <p>When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.</p> </div> <div data-bbox="488 1576 592 1771" style="background-color: black; color: white; padding: 5px;"> <p>NOTE</p> <p>If you are using Azure File storage, you cannot enable FIPS mode.</p> </div>	<p><code>false</code> or <code>true</code></p>

Parameter	Description	Values
<code>imageContentSources:</code>	Sources and repositories for the release-image content.	Array of objects. Includes a source and, optionally, mirrors , as described in the following rows of this table.
<code>imageContentSources: source:</code>	Required if you use imageContentSources . Specify the repository that users refer to, for example, in image pull specifications.	String
<code>imageContentSources: mirrors:</code>	Specify one or more repositories that may also contain the same images.	Array of strings
<code>publish:</code>	How to publish or expose the user-facing endpoints of your cluster, such as the Kubernetes API, OpenShift routes.	<p>Internal or External. The default value is External.</p> <p>Setting this field to Internal is not supported on non-cloud platforms.</p> <div data-bbox="970 1391 1078 1648" data-label="Image"> </div> <p>IMPORTANT</p> <p>If the value of the field is set to Internal, the cluster will become non-functional. For more information, refer to BZ#1953035.</p>

Parameter	Description	Values
sshKey:	<p>The SSH key to authenticate access to your cluster machines.</p>  <p>NOTE</p> <p>For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your ssh-agent process uses.</p>	<p>For example, sshKey: ssh-ed25519 AAAA...</p>

1. Not all CCO modes are supported for all cloud providers. For more information about CCO modes, see the "Managing cloud provider credentials" entry in the *Authentication and authorization* content.

21.10.1.4. Additional Red Hat OpenStack Platform (RHOSP) configuration parameters

Additional RHOSP configuration parameters are described in the following table:

Table 21.9. Additional RHOSP parameters

Parameter	Description	Values
compute: platform: openstack: rootVolume: size:	<p>For compute machines, the size in gigabytes of the root volume. If you do not set this value, machines use ephemeral storage.</p>	<p>Integer, for example 30.</p>
compute: platform: openstack: rootVolume: types:	<p>For compute machines, the root volume types.</p>	<p>A list of strings, for example, {performance-host1, performance-host2, performance-host3}.^[1]</p>

Parameter	Description	Values
compute: platform: openstack: rootVolume: type:	For compute machines, the root volume's type. This property is deprecated and is replaced by compute.platform.openstack.rootVolume.types .	String, for example, performance . ^[2]
compute: platform: openstack: rootVolume: zones:	For compute machines, the Cinder availability zone to install root volumes on. If you do not set a value for this parameter, the installation program selects the default availability zone. This parameter is mandatory when compute.platform.openstack.zones is defined.	A list of strings, for example ["zone-1", "zone-2"] .
controlPlane: platform: openstack: rootVolume: size:	For control plane machines, the size in gigabytes of the root volume. If you do not set this value, machines use ephemeral storage.	Integer, for example 30 .
controlPlane: platform: openstack: rootVolume: types:	For control plane machines, the root volume types.	A list of strings, for example, {performance-host1, performance-host2, performance-host3} . ^[1]
controlPlane: platform: openstack: rootVolume: type:	For control plane machines, the root volume's type. This property is deprecated and is replaced by compute.platform.openstack.rootVolume.types .	String, for example, performance . ^[2]

Parameter	Description	Values
controlPlane: platform: openstack: rootVolume: zones:	<p>For control plane machines, the Cinder availability zone to install root volumes on. If you do not set this value, the installation program selects the default availability zone. This parameter is mandatory when controlPlane.platform.openstack.zones is defined.</p>	A list of strings, for example ["zone-1", "zone-2"] .
platform: openstack: cloud:	<p>The name of the RHOSP cloud to use from the list of clouds in the clouds.yaml file.</p> <p>In the cloud configuration in the clouds.yaml file, if possible, use application credentials rather than a user name and password combination. Using application credentials avoids disruptions from secret propagation that follow user name and password rotation.</p>	String, for example MyCloud .
platform: openstack: externalNetwork:	<p>The RHOSP external network name to be used for installation.</p>	String, for example external .
platform: openstack: computeFlavor:	<p>The RHOSP flavor to use for control plane and compute machines.</p> <p>This property is deprecated. To use a flavor as the default for all machine pools, add it as the value of the type key in the platform.openstack.defaultMachinePlatform property. You can also set a flavor value for each machine pool individually.</p>	String, for example m1.xlarge .

1. If the machine pool defines **zones**, the count of types can either be a single item or match the number of items in **zones**. For example, the count of types cannot be 2 if there are 3 items in **zones**.
2. If you have any existing reference to this property, the installer populates the corresponding value in the **controlPlane.platform.openstack.rootVolume.types** field.

21.10.1.5. Optional RHOSP configuration parameters

Optional RHOSP configuration parameters are described in the following table:

Table 21.10. Optional RHOSP parameters

Parameter	Description	Values
<pre>compute: platform: openstack: additionalNetworkIDs:</pre>	Additional networks that are associated with compute machines. Allowed address pairs are not created for additional networks.	A list of one or more UUIDs as strings. For example, fa806b2f-ac49-4bce-b9db-124bc64209bf .
<pre>compute: platform: openstack: additionalSecurityGroupIDs:</pre>	Additional security groups that are associated with compute machines.	A list of one or more UUIDs as strings. For example, 7ee219f3-d2e9-48a1-96c2-e7429f1b0da7 .
<pre>compute: platform: openstack: zones:</pre>	RHOSP Compute (Nova) availability zones (AZs) to install machines on. If this parameter is not set, the installation program relies on the default settings for Nova that the RHOSP administrator configured.	A list of strings. For example, ["zone-1", "zone-2"] .

Parameter	Description	Values
<pre>compute: platform: openstack: serverGroupPolicy:</pre>	<p>Server group policy to apply to the group that will contain the compute machines in the pool. You cannot change server group policies or affiliations after creation. Supported options include anti-affinity, soft-affinity, and soft-anti-affinity. The default value is soft-anti-affinity.</p> <p>An affinity policy prevents migrations and therefore affects RHOSP upgrades. The affinity policy is not supported.</p> <p>If you use a strict anti-affinity policy, an additional RHOSP host is required during instance migration.</p>	<p>A server group policy to apply to the machine pool. For example, soft-affinity.</p>
<pre>controlPlane: platform: openstack: additionalNetworkIDs:</pre>	<p>Additional networks that are associated with control plane machines. Allowed address pairs are not created for additional networks.</p> <p>Additional networks that are attached to a control plane machine are also attached to the bootstrap node.</p>	<p>A list of one or more UUIDs as strings. For example, fa806b2f-ac49-4bce-b9db-124bc64209bf.</p>
<pre>controlPlane: platform: openstack: additionalSecurityGroupIDs:</pre>	<p>Additional security groups that are associated with control plane machines.</p>	<p>A list of one or more UUIDs as strings. For example, 7ee219f3-d2e9-48a1-96c2-e7429f1b0da7.</p>

Parameter	Description	Values
controlPlane: platform: openstack: zones:	RHOSP Compute (Nova) availability zones (AZs) to install machines on. If this parameter is not set, the installation program relies on the default settings for Nova that the RHOSP administrator configured.	A list of strings. For example, ["zone-1", "zone-2"].
controlPlane: platform: openstack: serverGroupPolicy:	<p>Server group policy to apply to the group that will contain the control plane machines in the pool. You cannot change server group policies or affiliations after creation. Supported options include anti-affinity, soft-affinity, and soft-anti-affinity. The default value is soft-anti-affinity.</p> <p>An affinity policy prevents migrations, and therefore affects RHOSP upgrades. The affinity policy is not supported.</p> <p>If you use a strict anti-affinity policy, an additional RHOSP host is required during instance migration.</p>	A server group policy to apply to the machine pool. For example, soft-affinity .
platform: openstack: clusterOSImage:	<p>The location from which the installation program downloads the RHCOS image.</p> <p>You must set this parameter to perform an installation in a restricted network.</p>	<p>An HTTP or HTTPS URL, optionally with an SHA-256 checksum.</p> <p>For example, http://mirror.example.com/images/rhcos-43.81.201912131630.0-openstack.x86_64.qcow2.gz?sha256=ffebbd68e8a1f2a245ca19522c16c86f67f9ac8e4e0c1f0a812b068b16f7265d. The value can also be the name of an existing Glance image, for example my-rhcos.</p>

Parameter	Description	Values
<pre>platform: openstack: clusterOSImageProperties:</pre>	<p>Properties to add to the installer-uploaded ClusterOSImage in Glance. This property is ignored if platform.openstack.clusterOSImage is set to an existing Glance image.</p> <p>You can use this property to exceed the default persistent volume (PV) limit for RHOSP of 26 PVs per node. To exceed the limit, set the hw_scsi_model property value to virtio-scsi and the hw_disk_bus value to scsi.</p> <p>You can also use this property to enable the QEMU guest agent by including the hw_qemu_guest_agent property with a value of yes.</p>	<p>A list of key-value string pairs. For example, ["hw_scsi_model": "virtio-scsi", "hw_disk_bus": "scsi"].</p>
<pre>platform: openstack: defaultMachinePlatform:</pre>	<p>The default machine pool platform configuration.</p>	<pre>{ "type": "ml.large", "rootVolume": { "size": 30, "type": "performance" } }</pre>
<pre>platform: openstack: ingressFloatingIP:</pre>	<p>An existing floating IP address to associate with the Ingress port. To use this property, you must also define the platform.openstack.externalNetwork property.</p>	<p>An IP address, for example 128.0.0.1.</p>
<pre>platform: openstack: apiFloatingIP:</pre>	<p>An existing floating IP address to associate with the API load balancer. To use this property, you must also define the platform.openstack.externalNetwork property.</p>	<p>An IP address, for example 128.0.0.1.</p>

Parameter	Description	Values
platform: openstack: externalDNS:	IP addresses for external DNS servers that cluster instances use for DNS resolution.	A list of IP addresses as strings. For example, ["8.8.8.8", "192.168.1.12"] .
platform: openstack: loadbalancer:	Whether or not to use the default, internal load balancer. If the value is set to UserManaged , this default load balancer is disabled so that you can deploy a cluster that uses an external, user-managed load balancer. If the parameter is not set, or if the value is OpenShiftManagedDefault , the cluster uses the default load balancer.	UserManaged or OpenShiftManagedDefault .
platform: openstack: machinesSubnet:	<p>The UUID of a RHOSP subnet that the cluster's nodes use. Nodes and virtual IP (VIP) ports are created on this subnet.</p> <p>The first item in networking.machineNetwork must match the value of machinesSubnet.</p> <p>If you deploy to a custom subnet, you cannot specify an external DNS server to the OpenShift Container Platform installer. Instead, add DNS to the subnet in RHOSP.</p>	A UUID as a string. For example, fa806b2f-ac49-4bce-b9db-124bc64209bf .

CHAPTER 22. INSTALLING ON OCI

22.1. INSTALLING A CLUSTER ON ORACLE CLOUD INFRASTRUCTURE (OCI) BY USING THE ASSISTED INSTALLER

From OpenShift Container Platform 4.16 and later versions, you can use the Assisted Installer to install a cluster on Oracle® Cloud Infrastructure (OCI) by using infrastructure that you provide.

22.1.1. The Assisted Installer and OCI overview

You can run cluster workloads on Oracle® Cloud Infrastructure (OCI) infrastructure that supports dedicated, hybrid, public, and multiple cloud environments. Both Red Hat and Oracle test, validate, and support running OCI in an OpenShift Container Platform cluster on OCI.

The Assisted Installer supports the OCI platform, and you can use the Assisted Installer to access an intuitive interactive workflow for the purposes of automating cluster installation tasks on OCI.

OCI provides services that can meet your needs for regulatory compliance, performance, and cost-effectiveness. You can access OCI Resource Manager configurations to provision and configure OCI resources.



IMPORTANT

The steps for provisioning OCI resources are provided as an example only. You can also choose to create the required resources through other methods; the scripts are just an example. Installing a cluster with infrastructure that you provide requires knowledge of the cloud provider and the installation process on OpenShift Container Platform. You can access OCI Resource Manager configurations to complete these steps, or use the configurations to model your own custom script.

Follow the steps in the *Installing a cluster on Oracle Cloud Infrastructure (OCI) by using the Assisted Installer* document to understand how to use the Assisted Installer to install a OpenShift Container Platform cluster on OCI. This document demonstrates the use of the OCI Cloud Controller Manager (CCM) and Oracle's Container Storage Interface (CSI) objects to link your OpenShift Container Platform cluster with the OCI API.



IMPORTANT

To ensure the best performance conditions for your cluster workloads that operate on OCI, ensure that volume performance units (VPUs) for your block volume are sized for your workloads. The following list provides guidance for selecting the VPUs needed for specific performance needs:

- Test or proof of concept environment: 100 GB, and 20 to 30 VPUs.
- Basic environment: 500 GB, and 60 VPUs.
- Heavy production environment: More than 500 GB, and 100 or more VPUs.

Consider reserving additional VPUs to provide sufficient capacity for updates and scaling activities. For more information about VPUs, see [Volume Performance Units \(Oracle documentation\)](#).

If you are unfamiliar with the OpenShift Container Platform Assisted Installer, see "Assisted Installer for OpenShift Container Platform".

Additional resources

- [Assisted Installer for OpenShift Container Platform](#)
- [Internet access for OpenShift Container Platform](#)
- [Volume Performance Units \(Oracle documentation\)](#)
- [Instance Sizing Recommendations for OpenShift Container Platform on OCI Nodes \(Oracle documentation\)](#)

22.1.2. Creating OCI resources and services

Create Oracle® Cloud Infrastructure (OCI) resources and services so that you can establish infrastructure with governance standards that meets your organization's requirements.

Prerequisites

- You configured an OCI account to host the cluster. See [Prerequisites \(Oracle documentation\)](#).

Procedure

1. Log in to your [Oracle Cloud Infrastructure \(OCI\)](#) account with administrator privileges.
2. Download an archive file from an Oracle resource. The archive file includes files for creating cluster resources and custom manifests. The archive file also includes a script, and when you run the script, the script creates OCI resources, such as DNS records, an instance, and so on. For more information, see [Configuration Files \(Oracle documentation\)](#).

22.1.3. Using the Assisted Installer to generate an OCI-compatible discovery ISO image

Generate a discovery ISO image and upload the image to Oracle® Cloud Infrastructure (OCI), so that the agent can perform hardware and network validation checks before you install an OpenShift Container Platform cluster on OCI.

From the OCI web console, you must create the following resources:

- A compartment for better organizing, restricting access, and setting usage limits to OCI resources.
- An object storage bucket for safely and securely storing the discovery ISO image. You can access the image at a later stage for the purposes of booting the instances, so that you can then create your cluster.

Prerequisites

- You created a child compartment and an object storage bucket on OCI. See [Provisioning Cloud Infrastructure \(OCI Console\)](#) in the Oracle documentation.
- You reviewed details about the OpenShift Container Platform installation and update processes.

- If you use a firewall and you plan to use a Telemetry service, you configured your firewall to allow OpenShift Container Platform to access the sites required.
- Before you create a virtual machines (VM), see [Cloud instance types \(Red Hat Ecosystem Catalog portal\)](#) to identify the supported OCI VM shapes.

Procedure

1. From the [Install OpenShift with the Assisted Installer](#) page on the Hybrid Cloud Console, generate the discovery ISO image by completing all the required Assisted Installer steps.
 - a. In the **Cluster Details** step, complete the following fields:

Field	Action required
Cluster name	Specify the name of your cluster, such as ocidemo .
Base domain	Specify the base domain of the cluster, such as splat-oci.devcluster.openshift.com . Provided you previously created a compartment on OCI, you can get this information by going to DNS management → Zones → List scope and then selecting the parent compartment. Your base domain should show under the Public zones tab.
OpenShift version	Specify OpenShift 4.16 or a later version.
CPU architecture	Specify x86_64 or Arm64 .
Integrate with external partner platforms	Specify Oracle Cloud Infrastructure . After you specify this value, the Include custom manifests checkbox is selected by default.

- b. On the **Operators** page, click **Next**.
- c. On the **Host Discovery** page, click **Add hosts**.
- d. For the **SSH public key** field, add your SSH key from your local system.

TIP

You can create an SSH authentication key pair by using the **ssh-keygen** tool.

- e. Click **Generate Discovery ISO** to generate the discovery ISO image file.
 - f. Download the file to your local system.
2. Upload the discovery ISO image to the OCI bucket. See [Uploading an Object Storage Object to a Bucket \(Oracle documentation\)](#).

- a. You must create a pre-authenticated request for your uploaded discovery ISO image. Ensure that you make note of the URL from the pre-authenticated request, because you must specify the URL at a later stage when you create an OCI stack.

Additional resources

- [Installation and update](#)
- [Configuring your firewall](#)

22.1.4. Provisioning OCI infrastructure for your cluster

By using the Assisted Installer to create details for your OpenShift Container Platform cluster, you can specify these details in a stack. A stack is an OCI feature where you can automate the provisioning of all necessary OCI infrastructure resources, such as the custom image, that are required for installing an OpenShift Container Platform cluster on OCI.

The Oracle® Cloud Infrastructure (OCI) Compute Service creates a virtual machine (VM) instance on OCI. This instance can then automatically attach to a virtual network interface controller (vNIC) in the virtual cloud network (VNC) subnet. On specifying the IP address of your OpenShift Container Platform cluster in the custom manifest template files, the OCI instance can communicate with your cluster over the VNC.

Prerequisites

- You uploaded the discovery ISO image to the OCI bucket. For more information, see "Using the Assisted Installer to generate an OCI-compatible discovery ISO image".

Procedure

1. Complete the steps for provisioning OCI infrastructure for your OpenShift Container Platform cluster. See [Creating OpenShift Container Platform Infrastructure Using Resource Manager \(Oracle documentation\)](#).
2. Create a stack, and then edit the custom manifest files according to the steps in the [Editing the OpenShift Custom Manifests \(Oracle documentation\)](#).

22.1.5. Completing the remaining Assisted Installer steps

After you provision Oracle® Cloud Infrastructure (OCI) resources and upload OpenShift Container Platform custom manifest configuration files to OCI, you must complete the remaining cluster installation steps on the Assisted Installer before you can create an instance OCI.

Prerequisites

- You created a resource stack on OCI that includes the custom manifest configuration files and OCI Resource Manager configuration resources. See "Provisioning OCI infrastructure for your cluster".

Procedure

1. From the [Red Hat Hybrid Cloud Console](#) web console, go to the **Host discovery** page.
2. Under the **Role** column, select either **Control plane node** or **Worker** for each targeted hostname.



IMPORTANT

Before, you can continue to the next steps, wait for each node to reach the **Ready** status.

3. Accept the default settings for the **Storage** and **Networking** steps, and then click **Next**.
4. On the **Custom manifests** page, in the **Folder** field, select **manifest**. This is the Assisted Installer folder where you want to save the custom manifest file.
 - a. In the **File name** field, enter a value such as **oci-ccm.yml**.
 - b. From the **Content** section, click **Browse**, and select the CCM manifest from your drive located in **custom_manifest/manifests/oci-ccm.yml**.
5. Expand the next **Custom manifest** section and repeat the same steps for the following manifests:
 - CSI driver manifest: **custom_manifest/manifests/oci-csi.yml**
 - CCM machine configuration: **custom_manifest/openshift/machineconfig-ccm.yml**
 - CSI driver machine configuration: **custom_manifest/openshift/machineconfig-csi.yml**
6. From the **Review and create** page, click **Install cluster** to create your OpenShift Container Platform cluster on OCI.

After the cluster installation and initialization operations, the Assisted Installer indicates the completion of the cluster installation operation. For more information, see "Completing the installation" section in the *Assisted Installer for OpenShift Container Platform* document.

Additional resources

- [Assisted Installer for OpenShift Container Platform](#)

22.1.6. Verifying a successful cluster installation on OCI

Verify that your cluster was installed and is running effectively on Oracle® Cloud Infrastructure (OCI).

Procedure

1. From the Hybrid Cloud Console, go to **Clusters > Assisted Clusters** and select your cluster's name.
2. Check that the Installation progress bar is at 100% and a message displays indicating "Installation completed successfully".
3. To access the OpenShift Container Platform web console, click the provided Web Console URL.
4. Go to the **Nodes** menu page.
5. Locate your node from the **Nodes** table.
6. From the **Overview** tab, check that your node has a **Ready** status.
7. Select the YAML tab.

- Check the **labels** parameter, and verify that the listed labels apply to your configuration. For example, the **topology.kubernetes.io/region=us-sanjose-1** label indicates in what OCI region the node was deployed.

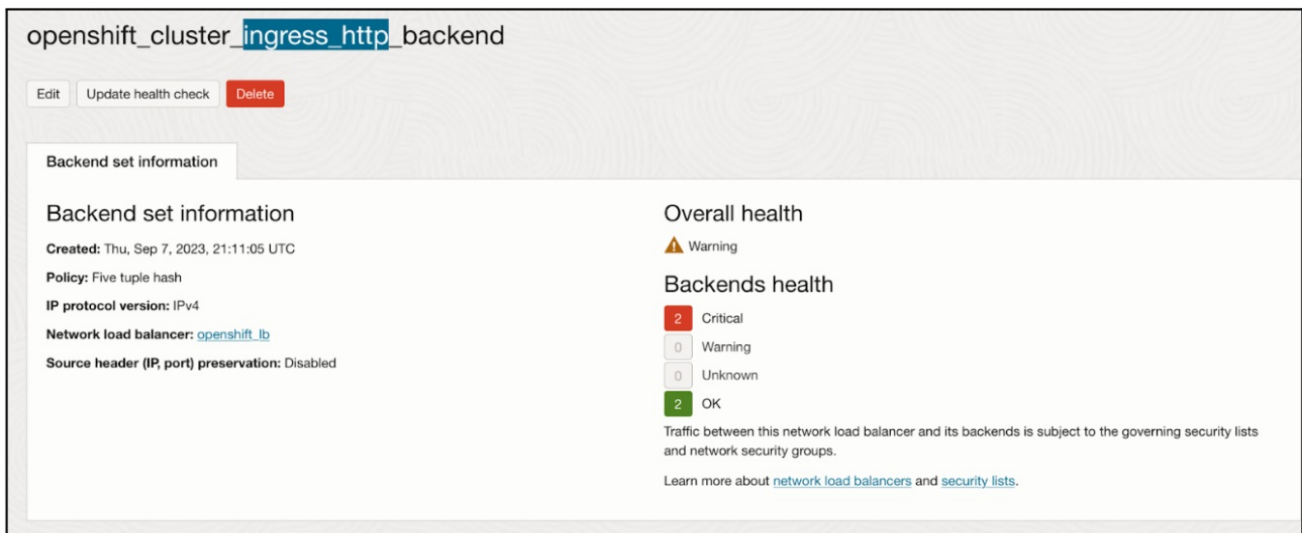
22.1.7. Troubleshooting the installation of a cluster on OCI

If you experience issues with using the Assisted Installer to install an OpenShift Container Platform cluster on Oracle® Cloud Infrastructure (OCI), read the following sections to troubleshoot common problems.

The Ingress Load Balancer in OCI is not at a healthy status

This issue is classed as a **Warning** because by using the Resource Manager to create a stack, you created a pool of compute nodes, 3 by default, that are automatically added as backend listeners for the Ingress Load Balancer. By default, the OpenShift Container Platform deploys 2 router pods, which are based on the default values from the OpenShift Container Platform manifest files. The **Warning** is expected because a mismatch exists with the number of router pods available, two, to run on the three compute nodes.

Figure 22.1. Example of a **Warning** message that is under the **Backend** set information tab on OCI:



You do not need to modify the Ingress Load Balancer configuration. Instead, you can point the Ingress Load Balancer to specific compute nodes that operate in your cluster on OpenShift Container Platform. To do this, use placement mechanisms, such as annotations, on OpenShift Container Platform to ensure router pods only run on the compute nodes that you originally configured on the Ingress Load Balancer as backend listeners.

OCI create stack operation fails with an Error: 400-InvalidParameter message

On attempting to create a stack on OCI, you identified that the **Logs** section of the job outputs an error message. For example:

Error: 400-InvalidParameter, DNS Label oci-demo does not follow Oracle requirements
 Suggestion: Please update the parameter(s) in the Terraform config as per error message DNS Label oci-demo does not follow Oracle requirements
 Documentation: https://registry.terraform.io/providers/oracle/oci/latest/docs/resources/core_vc

Go to the [Install OpenShift with the Assisted Installer](#) page on the Hybrid Cloud Console, and check the **Cluster name** field on the **Cluster Details** step. Remove any special characters, such as a hyphen (-), from the name, because these special characters are not compatible with the OCI naming conventions. For example, change **oci-demo** to **ocidemo**.

Additional resources

- [Troubleshooting OpenShift Container Platform on OCI \(Oracle documentation\)](#)
- [Installing an on-premise cluster using the Assisted Installer](#)

22.2. INSTALLING A CLUSTER ON ORACLE CLOUD INFRASTRUCTURE (OCI) BY USING THE AGENT-BASED INSTALLER

In OpenShift Container Platform 4.16, you can use the Agent-based Installer to install a cluster on Oracle® Cloud Infrastructure (OCI), so that you can run cluster workloads on infrastructure that supports dedicated, hybrid, public, and multiple cloud environments.

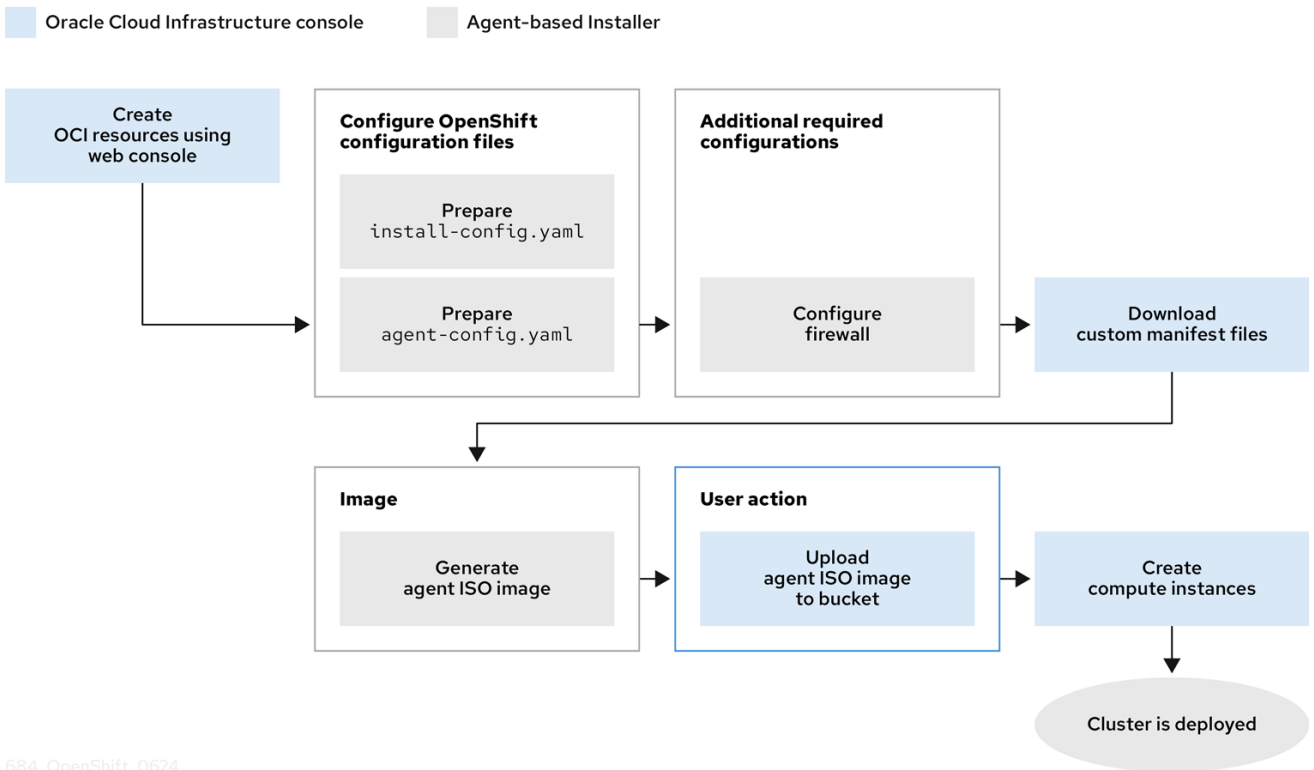
22.2.1. The Agent-based Installer and OCI overview

You can install an OpenShift Container Platform cluster on Oracle® Cloud Infrastructure (OCI) by using the Agent-based Installer. Both Red Hat and Oracle test, validate, and support running OCI and Oracle® Cloud VMware Solution (OCVS) workloads in an OpenShift Container Platform cluster on OCI.

The Agent-based installer provides the ease of use of the Assisted Installation service, but with the capability to install a cluster in either a connected or disconnected environment.

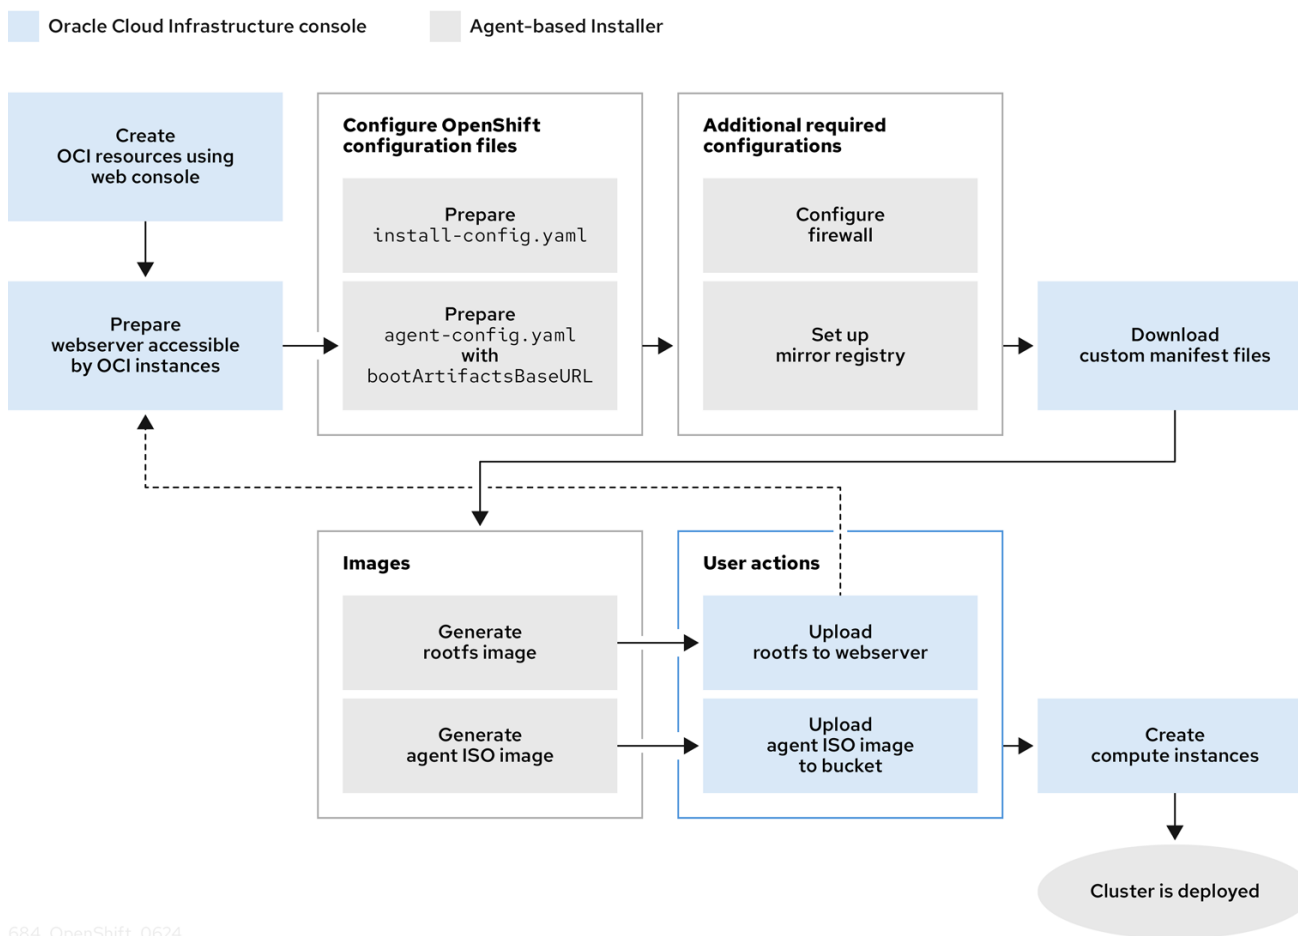
The following diagrams show workflows for connected and disconnected environments:

Figure 22.2. Workflow for using the Agent-based installer in a connected environment to install a cluster on OCI



684_OpenShift_0624

Figure 22.3. Workflow for using the Agent-based installer in a disconnected environment to install a cluster on OCI



684_OpenShift_0624

OCI provides services that can meet your regulatory compliance, performance, and cost-effectiveness needs. OCI supports 64-bit **x86** instances and 64-bit **ARM** instances. Additionally, OCI provides an OCVS service where you can move VMware workloads to OCI with minimal application re-architecture.




NOTE

Consider selecting a nonvolatile memory express (NVMe) drive or a solid-state drive (SSD) for your boot disk, because these drives offer low latency and high throughput capabilities for your boot disk.

By running your OpenShift Container Platform cluster on OCI, you can access the following capabilities:

- Compute flexible shapes, where you can customize the number of Oracle® CPUs (OCPU) and memory resources for your VM. With access to this capability, a cluster's workload can perform operations in a resource-balanced environment. You can find all RHEL-certified OCI shapes by going to the Oracle page on the Red Hat Ecosystem Catalog portal.
- Block Volume storage, where you can configure scaling and auto-tuning settings for your storage volume, so that the Block Volume service automatically adjusts the performance level to optimize performance.
- OCVS, where you can deploy a cluster in a public-cloud environment that operates on a VMware® vSphere software-defined data center (SDDC). You continue to retain full-administrative control over your VMware vSphere environment, but you can use OCI services to

improve your applications on flexible, scalable, and secure infrastructure.



IMPORTANT

To ensure the best performance conditions for your cluster workloads that operate on OCI and on the OCVS service, ensure volume performance units (VPUs) for your block volume is sized for your workloads. The following list provides some guidance in selecting the VPUs needed for specific performance needs:

- Test or proof of concept environment: 100 GB, and 20 to 30 VPUs.
- Basic environment: 500 GB, and 60 VPUs.
- Heavy production environment: More than 500 GB, and 100 or more VPUs.

Consider reserving additional VPUs to provide sufficient capacity for updates and scaling activities. For more information about VPUs, see [Volume Performance Units \(Oracle documentation\)](#).


Additional resources

- [Installation process](#)
- [Internet access for OpenShift Container Platform](#)
- [Understanding the Agent-based Installer](#)
- [Overview of the Compute Service \(Oracle documentation\)](#)
- [Volume Performance Units \(Oracle documentation\)](#)
- [Instance Sizing Recommendations for OpenShift Container Platform on OCI Nodes \(Oracle documentation\)](#)

22.2.2. Creating OCI infrastructure resources and services

You must create an OCI environment on your virtual machine (VM) shape. By creating this environment, you can install OpenShift Container Platform and deploy a cluster on an infrastructure that supports a wide range of cloud options and strong security policies. Having prior knowledge of OCI components can help you with understanding the concept of OCI resources and how you can configure them to meet your organizational needs.

The Agent-based installer method for installing an OpenShift Container Platform cluster on OCI requires that you manually create OCI resources and services.



IMPORTANT

To ensure compatibility with OpenShift Container Platform, you must set **A** as the record type for each DNS record and name records as follows:

- **api.<cluster_name>.<base_domain>**, which targets the **apiVIP** parameter of the API load balancer.
- **api-int.<cluster_name>.<base_domain>**, which targets the **apiVIP** parameter of the API load balancer.
- ***.apps.<cluster_name>.<base_domain>**, which targets the **ingressVIP** parameter of the Ingress load balancer.

The **api.*** and **api-int.*** DNS records relate to control plane machines, so you must ensure that all nodes in your installed OpenShift Container Platform cluster can access these DNS records.

Prerequisites

- You configured an OCI account to host the OpenShift Container Platform cluster. See [Prerequisites \(Oracle documentation\)](#).

Procedure

- Create the required OCI resources and services. See [OCI Resources Needed for Using the Agent-based Installer \(Oracle documentation\)](#).

Additional resources

- [Learn About Oracle Cloud Basics \(Oracle documentation\)](#)

22.2.3. Creating configuration files for installing a cluster on OCI

You need to create the **install-config.yaml** and the **agent-config.yaml** configuration files so that you can use the Agent-based Installer to generate a bootable ISO image. The Agent-based installation comprises a bootable ISO that has the Assisted discovery agent and the Assisted Service. Both of these components are required to perform the cluster installation, but the latter component runs on only one of the hosts.

At a later stage, you must follow the steps in the Oracle documentation for uploading your generated agent ISO image to Oracle's default Object Storage bucket, which is the initial step for integrating your OpenShift Container Platform cluster on Oracle® Cloud Infrastructure (OCI).



NOTE

You can also use the Agent-based Installer to generate or accept Zero Touch Provisioning (ZTP) custom resources.


Prerequisites

- You reviewed details about the OpenShift Container Platform installation and update processes.

- You read the documentation on selecting a cluster installation method and preparing the method for users.
- You have read the "Preparing to install with the Agent-based Installer" documentation.
- You downloaded the Agent-Based Installer and the command-line interface (CLI) from the Red Hat Hybrid Cloud Console.
- You have logged in to the OpenShift Container Platform with administrator privileges.

Procedure

1. For a disconnected environment, mirror the Mirror registry for Red Hat OpenShift to your local container image registry.



IMPORTANT

Check that your **openshift-install** binary version relates to your local image container registry and not a shared registry, such as Red Hat Quay.

```
$ ./openshift-install version
```

Example output for a shared registry binary

```
./openshift-install 4.16.0
built from commit ae7977b7d1ca908674a0d45c5c243c766fa4b2ca
release image registry.ci.openshift.org/origin/release:4.16ocp-
release@sha256:0da6316466d60a3a4535d5fed3589feb0391989982fba59d47d
4c729912d6363
release architecture amd64
```

2. Configure the **install-config.yaml** configuration file to meet the needs of your organization.

Example **install-config.yaml** configuration file that demonstrates setting an external platform

```
# install-config.yaml
apiVersion: v1
baseDomain: <base_domain> ❶
networking:
  clusterNetwork:
    - cidr: 10.128.0.0/14
      hostPrefix: 23
  network type: OVNKubernetes
  machineNetwork:
    - cidr: <ip_address_from_cidr> ❷
  serviceNetwork:
    - 172.30.0.0/16
compute:
  - architecture: amd64 ❸
    hyperthreading: Enabled
    name: worker
    replicas: 0
```


```

controlPlane:
  architecture: amd64 4
  hyperthreading: Enabled
  name: master
  replicas: 3
platform:
  external:
    platformName: oci 5
    cloudControllerManager: External
sshKey: <public_ssh_key> 6
pullSecret: '<pull_secret>' 7
# ...

```

- 1 The base domain of your cloud provider.
- 2 The IP address from the virtual cloud network (VCN) that the CIDR allocates to resources and components that operate on your network.
- 3 4 Depending on your infrastructure, you can select either **x86_64**, or **amd64**.
- 5 Set **OCI** as the external platform, so that OpenShift Container Platform can integrate with OCI.
- 6 Specify your SSH public key.
- 7 The pull secret that you need for authenticate purposes when downloading container images for OpenShift Container Platform components and services, such as Quay.io. See [Install OpenShift Container Platform 4](#) from the Red Hat Hybrid Cloud Console.

3. Create a directory on your local system named **openshift**.



IMPORTANT

Do not move the **install-config.yaml** and **agent-config.yaml** configuration files to the **openshift** directory.

4. Complete the steps in the "[Configuration Files](#)" section of the *Oracle* documentation to download Oracle Cloud Controller Manager (CCM) and Oracle Container Storage Interface (CSI) manifests as an archive file and save the archive file in your **openshift** directory. You need the Oracle CCM manifests for deploying the Oracle CCM during cluster installation so that OpenShift Container Platform can connect to the external OCI platform. You need the Oracle CSI custom manifests for deploying the Oracle CSI driver during cluster installation so that OpenShift Container Platform can claim required objects from OCI.
5. Access the custom manifest files that are provided in the "[Configuration Files](#)" section of the *Oracle* documentation.
 - a. Change the **oci-cloud-controller-manager** secret that is defined in the **oci-ccm.yml** configuration file to match your organization's region, compartment OCID, VCN OCID, and the subnet OCID from the load balancer.
6. Use the Agent-based Installer to generate a minimal ISO image, which excludes the rootfs image, by entering the following command in your OpenShift Container Platform CLI. You can use this image later in the process to boot all your cluster's nodes.

—

```
$ ./openshift-install agent create image --log-level debug
```

The command also completes the following actions:

- Creates a subdirectory, `./<installation_directory>/auth directory;`, and places **kubeadmin-password** and **kubeconfig** files in the subdirectory.
- Creates a **rendezvousIP** file based on the IP address that you specified in the **agent-config.yaml** configuration file.
- Optional: Any modifications you made to **agent-config.yaml** and **install-config.yaml** configuration files get imported to the Zero Touch Provisioning (ZTP) custom resources.



IMPORTANT

The Agent-based Installer uses Red Hat Enterprise Linux CoreOS (RHCOS). The rootfs image, which is mentioned in a later listed item, is required for booting, recovering, and repairing your operating system.

7. Configure the **agent-config.yaml** configuration file to meet your organization's requirements.

Example agent-config.yaml configuration file that sets values for an IPv4 formatted network.

```
apiVersion: v1alpha1
metadata:
  name: <cluster_name> 1
  namespace: <cluster_namespace> 2
rendezvousIP: <ip_address_from_CIDR> 3
bootArtifactsBaseURL: <server_URL> 4
# ...
```

- 1 The cluster name that you specified in your DNS record.
- 2 The namespace of your cluster on OpenShift Container Platform.
- 3 If you use IPv4 as the network IP address format, ensure that you set the **rendezvousIP** parameter to an IPv4 address that the VCN's Classless Inter-Domain Routing (CIDR) method allocates on your network. Also ensure that at least one instance from the pool of instances that you booted with the ISO matches the IP address value you set for **rendezvousIP**.
- 4 The URL of the server where you want to upload the rootfs image.

8. Apply one of the following two updates to your **agent-config.yaml** configuration file:

- For a disconnected network: After you run the command to generate a minimal ISO Image, the Agent-based installer saves the rootfs image into the `./<installation_directory>/boot-artifacts` directory on your local system. Use your preferred web server, such as any Hypertext Transfer Protocol daemon (**httpd**), to upload rootfs to the location stated in the **bootArtifactsBaseURL** parameter in the **agent-config.yaml** configuration file. For example, if the **bootArtifactsBaseURL** parameter states **http://192.168.122.20**, you would upload the generated rootfs image to this location, so that the Agent-based installer

can access the image from http://192.168.122.20/agent.x86_64-rootfs.img. After the Agent-based installer boots the minimal ISO for the external platform, the Agent-based Installer downloads the rootfs image from the http://192.168.122.20/agent.x86_64-rootfs.img location into the system memory.



NOTE

The Agent-based Installer also adds the value of the **bootArtifactsBaseURL** to the minimal ISO Image's configuration, so that when the Operator boots a cluster's node, the Agent-based Installer downloads the rootfs image into system memory.

- For a connected network: You do not need to specify the **bootArtifactsBaseURL** parameter in the **agent-config.yaml** configuration file. The default behavior of the Agent-based Installer reads the rootfs URL location from <https://rhcos.mirror.openshift.com>. After the Agent-based Installer boots the minimal ISO for the external platform, the Agent-based Installer then downloads the rootfs file into your system's memory from the default RHCOS URL.



IMPORTANT

Consider that the full ISO image, which is in excess of **1 GB**, includes the rootfs image. The image is larger than the minimal ISO Image, which is typically less than **150 MB**.

Additional resources

- [About OpenShift Container Platform installation](#)
- [Selecting a cluster installation type](#)
- [Preparing to install with the Agent-based Installer](#)
- [Downloading the Agent-based Installer](#)
- [Mirroring the OpenShift Container Platform image repository](#)
- [Optional: Using ZTP manifests](#)

22.2.4. Configuring your firewall for OpenShift Container Platform

Before you install OpenShift Container Platform, you must configure your firewall to grant access to the sites that OpenShift Container Platform requires. When using a firewall, make additional configurations to the firewall so that OpenShift Container Platform can access the sites that it requires to function.

For a disconnected environment, you must mirror content from both Red Hat and Oracle. This environment requires that you create firewall rules to expose your firewall to specific ports and registries.



NOTE

If your environment has a dedicated load balancer in front of your OpenShift Container Platform cluster, review the allowlists between your firewall and load balancer to prevent unwanted network restrictions to your cluster.

Procedure

1. Set the following registry URLs for your firewall's allowlist:

URL	Port	Function
registry.redhat.io	443	Provides core container images
access.redhat.com ^[1]	443	Hosts all the container images that are stored on the Red Hat Ecosystem Catalog, including core container images.
quay.io	443	Provides core container images
cdn.quay.io	443	Provides core container images
cdn01.quay.io	443	Provides core container images
cdn02.quay.io	443	Provides core container images
cdn03.quay.io	443	Provides core container images
cdn04.quay.io	443	Provides core container images
cdn05.quay.io	443	Provides core container images
cdn06.quay.io	443	Provides core container images
sso.redhat.com	443	The https://console.redhat.com site uses authentication from sso.redhat.com

1. In a firewall environment, ensure that the **access.redhat.com** resource is on the allowlist. This resource hosts a signature store that a container client requires for verifying images when pulling them from **registry.access.redhat.com**.

You can use the wildcards ***.quay.io** and ***.openshiftapps.com** instead of **cdn.quay.io** and **cdn0[1-3].quay.io** in your allowlist. When you add a site, such as **quay.io**, to your allowlist, do not add a wildcard entry, such as ***.quay.io**, to your denylist. In most cases, image registries use a content delivery network (CDN) to serve images. If a firewall blocks access, image downloads are denied when the initial download request redirects to a hostname such as **cdn01.quay.io**.

2. Set your firewall's allowlist to include any site that provides resources for a language or framework that your builds require.
3. If you do not disable Telemetry, you must grant access to the following URLs to access Red Hat Insights:

URL	Port	Function
cert-api.access.redhat.com	443	Required for Telemetry
api.access.redhat.com	443	Required for Telemetry
infogw.api.openshift.com	443	Required for Telemetry
console.redhat.com	443	Required for Telemetry and for insights-operator

4. Set your firewall's allowlist to include the following registry URLs:

URL	Port	Function
api.openshift.com	443	Required both for your cluster token and to check if updates are available for the cluster.
rhcosp.mirror.openshift.com	443	Required to download Red Hat Enterprise Linux CoreOS (RHCOS) images.

5. Set your firewall's allowlist to include the following external URLs. Each repository URL hosts OCI containers. Consider mirroring images to as few repositories as possible to reduce any performance issues.

URL	Port	Function
k8s.gcr.io	port	A Kubernetes registry that hosts container images for a community-based image registry. This image registry is hosted on a custom Google Container Registry (GCR) domain.
ghcr.io	port	A GitHub image registry where you can store and manage Open Container Initiative images. Requires an access token to publish, install, and delete private, internal, and public packages.
storage.googleapis.com	443	A source of release image signatures, although the Cluster Version Operator needs only a single functioning source.
registry.k8s.io	port	Replaces the k8s.gcr.io image registry because the k8s.gcr.io image registry does not support other platforms and vendors.

22.2.5. Running a cluster on OCI

To run a cluster on Oracle® Cloud Infrastructure (OCI), you must upload the generated agent ISO image to the default Object Storage bucket on OCI. Additionally, you must create a compute instance from the supplied base image, so that your OpenShift Container Platform and OCI can communicate with each other for the purposes of running the cluster on OCI.



NOTE

OCI supports the following OpenShift Container Platform cluster topologies:

- Installing an OpenShift Container Platform cluster on a single node.
- A highly available cluster that has a minimum of three control plane instances and two compute instances.
- A compact three-node cluster that has a minimum of three control plane instances.

Prerequisites

- You generated an agent ISO image. See the "Creating configuration files for installing a cluster on OCI" section.

Procedure

1. Upload the agent ISO image to Oracle's default Object Storage bucket and import the agent ISO image as a custom image to this bucket. Ensure you that you configure the custom image to boot in Unified Extensible Firmware Interface (UEFI) mode. For more information, see [Creating the OpenShift Container Platform ISO Image \(Oracle documentation\)](#) .
2. Create a compute instance from the supplied base image for your cluster topology. See [Creating the OpenShift Container Platform cluster on OCI \(Oracle documentation\)](#) .



IMPORTANT

Before you create the compute instance, check that you have enough memory and disk resources for your cluster. Additionally, ensure that at least one compute instance has the same IP address as the address stated under **rendezvousIP** in the **agent-config.yaml** file.

Additional resources

- [Recommended resources for topologies](#)
- [Instance Sizing Recommendations for OpenShift Container Platform on OCI Nodes \(Oracle documentation\)](#)
- [Troubleshooting OpenShift Container Platform on OCI \(Oracle documentation\)](#)

22.2.6. Verifying that your Agent-based cluster installation runs on OCI

Verify that your cluster was installed and is running effectively on Oracle® Cloud Infrastructure (OCI).

Prerequisites

- You created all the required OCI resources and services. See the "Creating OCI infrastructure resources and services" section.
- You created **install-config.yaml** and **agent-config.yaml** configuration files. See the "Creating configuration files for installing a cluster on OCI" section.
- You uploaded the agent ISO image to Oracle's default Object Storage bucket, and you created a compute instance on OCI. For more information, see "Running a cluster on OCI".

Procedure

After you deploy the compute instance on a self-managed node in your OpenShift Container Platform cluster, you can monitor the cluster's status by choosing one of the following options:

- From the OpenShift Container Platform CLI, enter the following command:

```
$ ./openshift-install agent wait-for install-complete --log-level debug
```

Check the status of the **rendezvous** host node that runs the bootstrap node. After the host reboots, the host forms part of the cluster.

- Use the **kubeconfig** API to check the status of various OpenShift Container Platform components. For the **KUBECONFIG** environment variable, set the relative path of the cluster's **kubeconfig** configuration file:

```
$ export KUBECONFIG=~/.kube/kubeconfig
```

Check the status of each of the cluster's self-managed nodes. CCM applies a label to each node to designate the node as running in a cluster on OCI.

```
$ oc get nodes -A
```

Output example

```
NAME                                STATUS ROLES          AGE VERSION
main-0.private.agenttest.oraclevcn.com Ready control-plane, master 7m v1.27.4+6eeca63
main-1.private.agenttest.oraclevcn.com Ready control-plane, master 15m v1.27.4+d7fa83f
main-2.private.agenttest.oraclevcn.com Ready control-plane, master 15m v1.27.4+d7fa83f
```

Check the status of each of the cluster's Operators, with the CCM Operator status being a good indicator that your cluster is running.

```
$ oc get co
```

Truncated output example

```
NAME          VERSION  AVAILABLE PROGRESSING  DEGRADED  SINCE
MESSAGE
authentication 4.16.0-0 True    False    False    6m18s
baremetal     4.16.0-0 True    False    False    2m42s
network       4.16.0-0 True    True     False    5m58s Progressing: ...
...
```

Additional resources

- [Gathering log data from a failed Agent-based installation](#)

CHAPTER 23. INSTALLING ON VSPHERE

23.1. INSTALLATION METHODS

You can install an OpenShift Container Platform cluster on vSphere using a variety of different installation methods. Each method has qualities that can make them more suitable for different use cases, such as installing a cluster in a disconnected environment or installing a cluster with minimal configuration and provisioning.

23.1.1. Assisted Installer

You can install OpenShift Container Platform with the [Assisted Installer](#). This method requires no setup for the installer and is ideal for connected environments like vSphere. Installing with the Assisted Installer also provides integration with vSphere, enabling autoscaling. See [Installing an on-premise cluster using the Assisted Installer](#) for additional details.

23.1.2. Agent-based Installer

You can install an OpenShift Container Platform cluster on vSphere using the Agent-based Installer. The Agent-based Installer can be used to boot an on-premises server in a disconnected environment by using a bootable image. With the Agent-based Installer, users also have the flexibility to provision infrastructure, customize network configurations, and customize installations within a disconnected environment. See [Preparing to install with the Agent-based Installer](#) for additional details.

23.1.3. Installer-provisioned infrastructure installation

You can install OpenShift Container Platform on vSphere by using installer-provisioned infrastructure. Installer-provisioned infrastructure allows the installation program to preconfigure and automate the provisioning of resources required by OpenShift Container Platform. Installer-provisioned infrastructure is useful for installing in environments with disconnected networks, where the installation program provisions the underlying infrastructure for the cluster.


- [Installing a cluster on vSphere](#) You can install OpenShift Container Platform on vSphere by using installer-provisioned infrastructure installation with no customization.
- [Installing a cluster on vSphere with customizations](#) You can install OpenShift Container Platform on vSphere by using installer-provisioned infrastructure installation with the default customization options.
- [Installing a cluster on vSphere with network customizations](#) You can install OpenShift Container Platform on installer-provisioned vSphere infrastructure, with network customizations. You can customize your OpenShift Container Platform network configuration during installation, so that your cluster can coexist with your existing IP address allocations and adhere to your network requirements.
- [Installing a cluster on vSphere in a restricted network](#) You can install a cluster on VMware vSphere infrastructure in a restricted network by creating an internal mirror of the installation release content. You can use this method to deploy OpenShift Container Platform on an internal network that is not visible to the internet.

23.1.4. User-provisioned infrastructure installation

You can install OpenShift Container Platform on vSphere by using user-provisioned infrastructure. User-provisioned infrastructure requires the user to provision all resources required by OpenShift

Container Platform. If you do not use infrastructure that the installation program provisions, you must manage and maintain the cluster resources yourself.

- [Installing a cluster on vSphere with user-provisioned infrastructure](#) You can install OpenShift Container Platform on VMware vSphere infrastructure that you provision.
- [Installing a cluster on vSphere with network customizations with user-provisioned infrastructure](#): You can install OpenShift Container Platform on VMware vSphere infrastructure that you provision with customized network configuration options.
- [Installing a cluster on vSphere in a restricted network with user-provisioned infrastructure](#) OpenShift Container Platform can be installed on VMware vSphere infrastructure that you provision in a restricted network.



IMPORTANT

The steps for performing a user-provisioned infrastructure installation are provided as an example only. Installing a cluster with infrastructure you provide requires knowledge of the vSphere platform and the installation process of OpenShift Container Platform. Use the user-provisioned infrastructure installation instructions as a guide; you are free to create the required resources through other methods.

23.1.5. Additional resources

- [Installation process](#)

23.2. INSTALLER-PROVISIONED INFRASTRUCTURE

23.2.1. vSphere installation requirements

Before you begin an installation using installer-provisioned infrastructure, be sure that your vSphere environment meets the following installation requirements.

23.2.1.1. VMware vSphere infrastructure requirements

You must install an OpenShift Container Platform cluster on one of the following versions of a VMware vSphere instance that meets the requirements for the components that you use:

- Version 7.0 Update 2 or later
- Version 8.0 Update 1 or later

Both of these releases support Container Storage Interface (CSI) migration, which is enabled by default on OpenShift Container Platform 4.16.

You can host the VMware vSphere infrastructure on-premise or on a [VMware Cloud Verified provider](#) that meets the requirements outlined in the following tables:

Table 23.1. Version requirements for vSphere virtual environments

Virtual environment product	Required version
VMware virtual hardware	15 or later

Virtual environment product	Required version
vSphere ESXi hosts	7.0 Update 2 or later; 8.0 Update 1 or later
vCenter host	7.0 Update 2 or later; 8.0 Update 1 or later




IMPORTANT

You must ensure that the time on your ESXi hosts is synchronized before you install OpenShift Container Platform. See [Edit Time Configuration for a Host](#) in the VMware documentation.

Table 23.2. Minimum supported vSphere version for VMware components

Component	Minimum supported versions	Description
Hypervisor	vSphere 7.0 Update 2 (or later) or vSphere 8.0 Update 1 (or later) with virtual hardware version 15	This hypervisor version is the minimum version that Red Hat Enterprise Linux CoreOS (RHCOS) supports. For more information about supported hardware on the latest version of Red Hat Enterprise Linux (RHEL) that is compatible with RHCOS, see Hardware on the Red Hat Customer Portal.
Optional: Networking (NSX-T)	vSphere 7.0 Update 2 or later; vSphere 8.0 Update 1 or later	At a minimum, vSphere 7.0 Update 2 or vSphere 8.0 Update 1 is required for OpenShift Container Platform. For more information about the compatibility of NSX and OpenShift Container Platform, see the Release Notes section of VMware's NSX container plugin documentation .
CPU micro-architecture	x86-64-v2 or higher	OpenShift 4.13 and later are based on RHEL 9.2 host operating system which raised the microarchitecture requirements to x86-64-v2. See the RHEL Microarchitecture requirements documentation . You can verify compatibility by following the procedures outlined in this KCS article .



IMPORTANT

To ensure the best performance conditions for your cluster workloads that operate on Oracle® Cloud Infrastructure (OCI) and on the Oracle® Cloud VMware Solution (OCVS) service, ensure volume performance units (VPUs) for your block volume are sized for your workloads.

The following list provides some guidance in selecting the VPUs needed for specific performance needs:

- Test or proof of concept environment: 100 GB, and 20 to 30 VPUs.
- Base-production environment: 500 GB, and 60 VPUs.
- Heavy-use production environment: More than 500 GB, and 100 or more VPUs.

Consider allocating additional VPUs to give enough capacity for updates and scaling activities. See [Block Volume Performance Levels \(Oracle documentation\)](#).

23.2.1.2. Network connectivity requirements

You must configure the network connectivity between machines to allow OpenShift Container Platform cluster components to communicate.

Review the following details about the required network ports.

Table 23.3. Ports used for all-machine to all-machine communications

Protocol	Port	Description
VRRP	N/A	Required for keepalived
ICMP	N/A	Network reachability tests
TCP	1936	Metrics
	9000-9999	Host level services, including the node exporter on ports 9100-9101 and the Cluster Version Operator on port 9099 .
	10250-10259	The default ports that Kubernetes reserves
	10256	openshift-sdn
UDP	4789	virtual extensible LAN (VXLAN)
	6081	Geneve
	9000-9999	Host level services, including the node exporter on ports 9100-9101 .
	500	IPsec IKE packets

Protocol	Port	Description
	4500	IPsec NAT-T packets
TCP/UDP	30000-32767	Kubernetes node port
ESP	N/A	IPsec Encapsulating Security Payload (ESP)

Table 23.4. Ports used for all-machine to control plane communications

Protocol	Port	Description
TCP	6443	Kubernetes API

Table 23.5. Ports used for control plane machine to control plane machine communications


Protocol	Port	Description
TCP	2379-2380	etcd server and peer ports

23.2.1.3. VMware vSphere CSI Driver Operator requirements

To install the vSphere Container Storage Interface (CSI) Driver Operator, the following requirements must be met:

- VMware vSphere version: 7.0 Update 2 or later; 8.0 Update 1 or later
- vCenter version: 7.0 Update 2 or later; 8.0 Update 1 or later
- Virtual machines of hardware version 15 or later
- No third-party vSphere CSI driver already installed in the cluster

If a third-party vSphere CSI driver is present in the cluster, OpenShift Container Platform does not overwrite it. The presence of a third-party vSphere CSI driver prevents OpenShift Container Platform from updating to OpenShift Container Platform 4.13 or later.



NOTE

The VMware vSphere CSI Driver Operator is supported only on clusters deployed with **platform: vsphere** in the installation manifest.

You can create a custom role for the Container Storage Interface (CSI) driver, the vSphere CSI Driver Operator, and the vSphere Problem Detector Operator. The custom role can include privilege sets that assign a minimum set of permissions to each vSphere object. This means that the CSI driver, the vSphere CSI Driver Operator, and the vSphere Problem Detector Operator can establish a basic interaction with these objects.



IMPORTANT

Installing an OpenShift Container Platform cluster in a vCenter is tested against a full list of privileges as described in the "Required vCenter account privileges" section. By adhering to the full list of privileges, you can reduce the possibility of unexpected and unsupported behaviors that might occur when creating a custom role with a set of restricted privileges.

Additional resources

- To remove a third-party vSphere CSI driver, see [Removing a third-party vSphere CSI Driver](#) .
- To update the hardware version for your vSphere nodes, see [Updating hardware on nodes running in vSphere](#).
- [Minimum permissions for the storage components](#)

23.2.1.4. vCenter requirements

Before you install an OpenShift Container Platform cluster on your vCenter that uses infrastructure that the installer provisions, you must prepare your environment.

Required vCenter account privileges

To install an OpenShift Container Platform cluster in a vCenter, the installation program requires access to an account with privileges to read and create the required resources. Using an account that has global administrative privileges is the simplest way to access all of the necessary permissions.

If you cannot use an account with global administrative privileges, you must create roles to grant the privileges necessary for OpenShift Container Platform cluster installation. While most of the privileges are always required, some are required only if you plan for the installation program to provision a folder to contain the OpenShift Container Platform cluster on your vCenter instance, which is the default behavior. You must create or amend vSphere roles for the specified objects to grant the required privileges.

An additional role is required if the installation program is to create a vSphere virtual machine folder.

Example 23.1. Roles and privileges required for installation in vSphere API

vSphere object for role	When required	Required privileges in vSphere API
-------------------------	---------------	------------------------------------

vSphere object for role	When required	Required privileges in vSphere API
vSphere vCenter	Always	Cns.Searchable InventoryService.Tagging.AttachTag InventoryService.Tagging.CreateCategory InventoryService.Tagging.CreateTag InventoryService.Tagging.DeleteCategory InventoryService.Tagging.DeleteTag InventoryService.Tagging.EditCategory InventoryService.Tagging.EditTag Sessions.ValidateSession StorageProfile.Update StorageProfile.View
vSphere vCenter Cluster	If VMs will be created in the cluster root	Host.Config.Storage Resource.AssignVMToPool VApp.AssignResourcePool VApp.Import VirtualMachine.Config.AddNewDisk
vSphere vCenter Resource Pool	If an existing resource pool is provided	Resource.AssignVMToPool VApp.AssignResourcePool VApp.Import VirtualMachine.Config.AddNewDisk
vSphere Datastore	Always	Datastore.AllocateSpace Datastore.Browse Datastore.FileManagement InventoryService.Tagging.ObjectAttachable
vSphere Port Group	Always	Network.Assign
Virtual Machine Folder	Always	InventoryService.Tagging.ObjectAttachable Resource.AssignVMToPool VApp.Import VirtualMachine.Config.AddExistingDisk VirtualMachine.Config.AddNewDisk VirtualMachine.Config.Add

vSphere object for role	When required	RemoveDevice Required privileges in vSphere API
		VirtualMachine.Config.AdvancedConfig VirtualMachine.Config.Annotation VirtualMachine.Config.CPUCount VirtualMachine.Config.DiskExtend VirtualMachine.Config.DiskLease VirtualMachine.Config.EditDevice VirtualMachine.Config.Memory VirtualMachine.Config.RemoveDisk VirtualMachine.Config.Rename Host.Config.Storage VirtualMachine.Config.ResetGuestInfo VirtualMachine.Config.Resource VirtualMachine.Config.Settings VirtualMachine.Config.UpgradeVirtualHardware VirtualMachine.Interact.GuestControl VirtualMachine.Interact.PowerOff VirtualMachine.Interact.PowerOn VirtualMachine.Interact.Reset VirtualMachine.Inventory.Create VirtualMachine.Inventory.CreateFromExisting VirtualMachine.Inventory.Delete VirtualMachine.Provisioning.Clone VirtualMachine.Provisioning.MarkAsTemplate VirtualMachine.Provisioning.DeployTemplate
vSphere vCenter Datacenter	If the installation program creates the virtual machine folder. For user-provisioned infrastructure, VirtualMachine.Inventory.Create and VirtualMachine.Inventory.D	InventoryService.Tagging.ObjectAttachable Resource.AssignVMToPool VirtualMachine.Config.AddExistingDisk VirtualMachine.Config.AddNewDisk

vSphere object for role	Required privileges are optional if your cluster does not use the Machine API. See the "Minimum permissions for the Machine API" table.	Required privileges in vSphere API
		VirtualMachine.Config.AddRemoveDevice VirtualMachine.Config.AdvancedConfig VirtualMachine.Config.Annotation VirtualMachine.Config.CPUCount VirtualMachine.Config.DiskExtend VirtualMachine.Config.DiskLease VirtualMachine.Config.EditDevice VirtualMachine.Config.Memory VirtualMachine.Config.RemoveDisk VirtualMachine.Config.Rename VirtualMachine.Config.ResetGuestInfo VirtualMachine.Config.Resource VirtualMachine.Config.Settings VirtualMachine.Config.UpgradeVirtualHardware VirtualMachine.Interact.GuestControl VirtualMachine.Interact.PowerOff VirtualMachine.Interact.PowerOn VirtualMachine.Interact.Reset VirtualMachine.Inventory.Create VirtualMachine.Inventory.CreateFromExisting VirtualMachine.Inventory.Delete VirtualMachine.Provisioning.Clone VirtualMachine.Provisioning.DeployTemplate VirtualMachine.Provisioning.MarkAsTemplate Folder.Create Folder.Delete

Example 23.2. Roles and privileges required for installation in vCenter graphical user interface (GUI)

vSphere object for role	When required	Required privileges in vCenter GUI
vSphere vCenter	Always	Cns.Searchable "vSphere Tagging"."Assign or Unassign vSphere Tag" "vSphere Tagging"."Create vSphere Tag Category" "vSphere Tagging"."Create vSphere Tag" vSphere Tagging"."Delete vSphere Tag Category" "vSphere Tagging"."Delete vSphere Tag" "vSphere Tagging"."Edit vSphere Tag Category" "vSphere Tagging"."Edit vSphere Tag" Sessions."Validate session" "Profile-driven storage"."Profile-driven storage update" "Profile-driven storage"."Profile-driven storage view"
vSphere vCenter Cluster	If VMs will be created in the cluster root	Host.Configuration."Storage partition configuration" Resource."Assign virtual machine to resource pool" VApp."Assign resource pool" VApp.Import "Virtual machine"."Change Configuration"."Add new disk"
vSphere vCenter Resource Pool	If an existing resource pool is provided	Host.Configuration."Storage partition configuration" Resource."Assign virtual machine to resource pool" VApp."Assign resource pool" VApp.Import "Virtual machine"."Change Configuration"."Add new disk"
vSphere Datastore	Always	Datastore."Allocate space" Datastore."Browse datastore" Datastore."Low level file operations" "vSphere Tagging"."Assign or Unassign vSphere Tag on Object"
vSphere Port Group	Always	Network."Assign network"

Virtual Machine Folder vSphere object for role	Always When required	"vSphere Tagging". "Assign or Unassign vSphere Tag on Object" Resource. "Assign virtual machine to resource pool" VApp.Import "Virtual machine". "Change Configuration". "Add existing disk" "Virtual machine". "Change Configuration". "Add new disk" "Virtual machine". "Change Configuration". "Add or remove device" "Virtual machine". "Change Configuration". "Advanced configuration" "Virtual machine". "Change Configuration". "Set annotation" "Virtual machine". "Change Configuration". "Change CPU count" "Virtual machine". "Change Configuration". "Extend virtual disk" "Virtual machine". "Change Configuration". "Acquire disk lease" "Virtual machine". "Change Configuration". "Modify device settings" "Virtual machine". "Change Configuration". "Change Memory" "Virtual machine". "Change Configuration". "Remove disk" "Virtual machine". "Change Configuration". Rename "Virtual machine". "Change Configuration". "Reset guest information" "Virtual machine". "Change Configuration". "Change resource" "Virtual machine". "Change Configuration". "Change Settings" "Virtual machine". "Change Configuration". "Upgrade virtual machine compatibility" "Virtual machine". Interaction. "Guest operating system management by VIX API" "Virtual machine". Interaction. "Power

vSphere object for role	When required	Required privileges in vCenter GUI
		<p>off" Virtual machine".Interaction."Power on" "Virtual machine".Interaction.Reset "Virtual machine"."Edit Inventory"."Create new" "Virtual machine"."Edit Inventory"."Create from existing" "Virtual machine"."Edit Inventory"."Remove" "Virtual machine".Provisioning."Clone virtual machine" "Virtual machine".Provisioning."Mark as template" "Virtual machine".Provisioning."Deploy template"</p>
vSphere vCenter Datacenter	<p>If the installation program creates the virtual machine folder. For user-provisioned infrastructure, VirtualMachine.Inventory.Create and VirtualMachine.Inventory.Delete privileges are optional if your cluster does not use the Machine API.</p>	<p>"vSphere Tagging"."Assign or Unassign vSphere Tag on Object" Resource."Assign virtual machine to resource pool" VApp.Import "Virtual machine"."Change Configuration"."Add existing disk" "Virtual machine"."Change Configuration"."Add new disk" "Virtual machine"."Change Configuration"."Add or remove device" "Virtual machine"."Change Configuration"."Advanced configuration" "Virtual machine"."Change Configuration"."Set annotation" "Virtual machine"."Change Configuration"."Change CPU count" "Virtual machine"."Change Configuration"."Extend virtual disk" "Virtual machine"."Change Configuration"."Acquire disk lease" "Virtual machine"."Change Configuration"."Modify device settings" "Virtual machine"."Change</p>

vSphere object for role	When required	Configuration". "Change Required privileges in vCenter GUI Memory "Virtual machine". "Change Configuration". "Remove disk" "Virtual machine". "Change Configuration". Rename "Virtual machine". "Change Configuration". "Reset guest information" "Virtual machine". "Change Configuration". "Change resource" "Virtual machine". "Change Configuration". "Change Settings" "Virtual machine". "Change Configuration". "Upgrade virtual machine compatibility" "Virtual machine". Interaction. "Guest operating system management by VIX API" "Virtual machine". Interaction. "Power off" "Virtual machine". Interaction. "Power on" "Virtual machine". Interaction. Reset "Virtual machine". "Edit Inventory". "Create new" "Virtual machine". "Edit Inventory". "Create from existing" "Virtual machine". "Edit Inventory". "Remove" "Virtual machine". Provisioning. "Clone virtual machine" "Virtual machine". Provisioning. "Deploy template" "Virtual machine". Provisioning. "Mark as template" Folder. "Create folder" Folder. "Delete folder"

Additionally, the user requires some **ReadOnly** permissions, and some of the roles require permission to propagate the permissions to child objects. These settings vary depending on whether or not you install the cluster into an existing folder.

Example 23.3. Required permissions and propagation settings

vSphere object	When required	Propagate to children	Permissions required
vSphere vCenter	Always	False	Listed required privileges
vSphere vCenter Datacenter	Existing folder	False	ReadOnly permission
	Installation program creates the folder	True	Listed required privileges
vSphere vCenter Cluster	Existing resource pool	False	ReadOnly permission
	VMs in cluster root	True	Listed required privileges
vSphere vCenter Datastore	Always	False	Listed required privileges
vSphere Switch	Always	False	ReadOnly permission
vSphere Port Group	Always	False	Listed required privileges
vSphere vCenter Virtual Machine Folder	Existing folder	True	Listed required privileges
vSphere vCenter Resource Pool	Existing resource pool	True	Listed required privileges

For more information about creating an account with only the required privileges, see [vSphere Permissions and User Management Tasks](#) in the vSphere documentation.

Minimum required vCenter account privileges

After you create a custom role and assign privileges to it, you can create permissions by selecting specific vSphere objects and then assigning the custom role to a user or group for each object.


Before you create permissions or request for the creation of permissions for a vSphere object, determine what minimum permissions apply to the vSphere object. By doing this task, you can ensure a basic interaction exists between a vSphere object and OpenShift Container Platform architecture.



IMPORTANT

If you create a custom role and you do not assign privileges to it, the vSphere Server by default assigns a **Read Only** role to the custom role. Note that for the cloud provider API, the custom role only needs to inherit the privileges of the **Read Only** role.

Consider creating a custom role when an account with global administrative privileges does not meet your needs.



IMPORTANT

Accounts that are not configured with the required privileges are unsupported. Installing an OpenShift Container Platform cluster in a vCenter is tested against a full list of privileges as described in the "Required vCenter account privileges" section. By adhering to the full list of privileges, you can reduce the possibility of unexpected behaviors that might occur when creating a custom role with a restricted set of privileges.

The following tables list the minimum permissions for a vSphere object that interacts with specific OpenShift Container Platform architecture.

Example 23.4. Minimum permissions on installer-provisioned infrastructure

vSphere object for role	When required	Required privileges
vSphere vCenter	Always	Cns.Searchable InventoryService.Tagging.AttachTag InventoryService.Tagging.CreateCategory InventoryService.Tagging.CreateTag InventoryService.Tagging.DeleteCategory InventoryService.Tagging.DeleteTag InventoryService.Tagging.EditCategory InventoryService.Tagging.EditTag Sessions.ValidateSession StorageProfile.Update StorageProfile.View
vSphere vCenter Cluster	If you intend to create VMs in the cluster root	Host.Config.StorageResource.AssignVMT oPool VApp.AssignResourcePool VApp.Import VirtualMachine.Config .AddNewDisk

vSphere object for role	When required	Required privileges
vSphere vCenter Resource Pool	If you provide an existing resource pool in the install-config.yaml file	Datastore.Browse Datastore.FileManagement Host.Config.Storage InventoryService.Tagging.ObjectAttachableResource.AssignVMT oPool VApp.AssignResourcePool VApp.Import`minimum
vSphere Port Group	Always	Network.Assign
Virtual Machine Folder	Always	InventoryService.Tagging.ObjectAttachableResource.AssignVMT oPool VApp.Import VirtualMachine.Config.AddExistingDisk VirtualMachine.Config.AddNewDisk VirtualMachine.Config.AddRemoveDevice VirtualMachine.Config.AdvancedConfig VirtualMachine.Config.Annotation VirtualMachine.Config.CPUCount VirtualMachine.Config.DiskExtend VirtualMachine.Config.DiskLease VirtualMachine.Config.EditDevice VirtualMachine.Config.Memory VirtualMachine.Config.RemoveDisk VirtualMachine.Config.Rename VirtualMachine.Config.ResetGuestInfo VirtualMachine.Config.Resource VirtualMachine.Config.Settings VirtualMachine.Config

vSphere object for role	When required	Required privileges
		UpgradeVirtualHardware VirtualMachine.Interact.GuestControl VirtualMachine.Interact.PowerOff VirtualMachine.Interact.PowerOn VirtualMachine.Interact.Reset VirtualMachine.Inventory.Create VirtualMachine.Inventory.CreateFromExisting VirtualMachine.Inventory.Delete VirtualMachine.Provisioning.Clone VirtualMachine.Provisioning.MarkAsTemplate VirtualMachine.Provisioning.DeployTemplate
vSphere vCenter Datacenter	<p>If the installation program creates the virtual machine folder. For user-provisioned infrastructure, VirtualMachine.Inventory.Create and VirtualMachine.Inventory.Delete privileges are optional if your cluster does not use the Machine API. If your cluster does use the Machine API and you want to set the minimum set of permissions for the API, see the "Minimum permissions for the Machine API" table.</p>	Folder.Create Folder.Delete InventoryService.Tagging.ObjectAttachableResource.AssignVMToolPool VApp.Import VirtualMachine.Config.AddExistingDisk VirtualMachine.Config.AddNewDisk VirtualMachine.Config.AddRemoveDevice VirtualMachine.Config.AdvancedConfig VirtualMachine.Config.Annotation VirtualMachine.Config.CPUCount VirtualMachine.Config.DiskExtend VirtualMachine.Config.DiskLease VirtualMachine.Config.EditDevice VirtualMachine.Config.Memory VirtualMachine.Config.RemoveDisk VirtualMachine.Config

vSphere object for role	When required	Required privileges
		.Rename VirtualMachine.Config .ResetGuestInfo VirtualMachine.Config .Resource VirtualMachine.Config .Settings VirtualMachine.Config .UpgradeVirtualHardware VirtualMachine.Interact.GuestControl VirtualMachine.Interact.PowerOff VirtualMachine.Interact.PowerOn VirtualMachine.Interact.Reset VirtualMachine.Inventory.Create VirtualMachine.Inventory.CreateFromExisting VirtualMachine.Inventory.Delete VirtualMachine.Provisioning.Clone VirtualMachine.Provisioning.DeployTemplate VirtualMachine.Provisioning.MarkAsTemplate

Example 23.5. Minimum permissions for post-installation management of components

vSphere object for role	When required	Required privileges
-------------------------	---------------	---------------------

vSphere object for role	When required	Required privileges
vSphere vCenter	Always	Cns.Searchable InventoryService.Tagging.AttachTag InventoryService.Tagging.CreateCategory InventoryService.Tagging.CreateTag InventoryService.Tagging.DeleteCategory InventoryService.Tagging.DeleteTag InventoryService.Tagging.EditCategory InventoryService.Tagging.EditTag Sessions.ValidateSession StorageProfile.Update StorageProfile.View
vSphere vCenter Cluster	If you intend to create VMs in the cluster root	Host.Config.StorageResource.AssignVMT oPool
vSphere vCenter Resource Pool	If you provide an existing resource pool in the install-config.yaml file	Host.Config.Storage
vSphere Datastore	Always	Datastore.AllocateSpace Datastore.Browse Datastore.FileManagement InventoryService.Tagging.ObjectAttachable
vSphere Port Group	Always	Network.Assign

vSphere object for role	When required	Required privileges
Virtual Machine Folder	Always	VirtualMachine.Config .AddExistingDisk VirtualMachine.Config .AddRemoveDevice VirtualMachine.Config .AdvancedConfig VirtualMachine.Config .Annotation VirtualMachine.Config .CPUCount VirtualMachine.Config .DiskExtend VirtualMachine.Config .Memory VirtualMachine.Config .Settings VirtualMachine.Interact t.PowerOff VirtualMachine.Interact t.PowerOn VirtualMachine.Inventory ory.CreateFromExisting VirtualMachine.Inventory ory.Delete VirtualMachine.Provisioning .Clone VirtualMachine.Provisioning .DeployTemplate
vSphere vCenter Datacenter	If the installation program creates the virtual machine folder. For user-provisioned infrastructure, VirtualMachine.Inventory.Create and VirtualMachine.Inventory.Delete privileges are optional if your cluster does not use the Machine API. If your cluster does use the Machine API and you want to set the minimum set of permissions for the API, see the "Minimum permissions for the Machine API" table.	Resource.AssignVMT oPool VirtualMachine.Config .AddExistingDisk VirtualMachine.Config .AddRemoveDevice VirtualMachine.Interact t.PowerOff VirtualMachine.Interact t.PowerOn VirtualMachine.Provisioning .DeployTemplate

Example 23.6. Minimum permissions for the storage components

vSphere object for role	When required	Required privileges
vSphere vCenter	Always	Cns.Searchable InventoryService.Tagging.CreateCategory InventoryService.Tagging.CreateTag InventoryService.Tagging.EditCategory InventoryService.Tagging.EditTag StorageProfile.Update StorageProfile.View
vSphere vCenter Cluster	If you intend to create VMs in the cluster root	Host.Config.Storage
vSphere vCenter Resource Pool	If you provide an existing resource pool in the install-config.yaml file	Host.Config.Storage
vSphere Datastore	Always	Datastore.Browse Datastore.FileManagement InventoryService.Tagging.ObjectAttachable
vSphere Port Group	Always	Read Only
Virtual Machine Folder	Always	VirtualMachine.Config.AddExistingDisk VirtualMachine.Config.AddRemoveDevice
vSphere vCenter Datacenter	If the installation program creates the virtual machine folder. For user-provisioned infrastructure, VirtualMachine.Inventory.Create and VirtualMachine.Inventory.Delete privileges are optional if your cluster does not use the Machine API. If your cluster does use the Machine API and you want to set the minimum set of permissions for the API, see the "Minimum permissions for the Machine API" table.	VirtualMachine.Config.AddExistingDisk VirtualMachine.Config.AddRemoveDevice

Example 23.7. Minimum permissions for the Machine API

vSphere object for role	When required	Required privileges
vSphere vCenter	Always	InventoryService.Tagging.AttachTag InventoryService.Tagging.CreateCategory InventoryService.Tagging.CreateTag InventoryService.Tagging.DeleteCategory InventoryService.Tagging.DeleteTag InventoryService.Tagging.EditCategory InventoryService.Tagging.EditTag Sessions.ValidateSession StorageProfile.Update StorageProfile.View
vSphere vCenter Cluster	If you intend to create VMs in the cluster root	Resource.AssignVMT oPool
vSphere vCenter Resource Pool	If you provide an existing resource pool in the install-config.yaml file	Read Only
vSphere Datastore	Always	Datastore.AllocateSpace Datastore.Browse
vSphere Port Group	Always	Network.Assign

vSphere object for role	When required	Required privileges
Virtual Machine Folder	Always	VirtualMachine.Config .AddRemoveDevice VirtualMachine.Config .AdvancedConfig VirtualMachine.Config .Annotation VirtualMachine.Config .CPUCount VirtualMachine.Config .DiskExtend VirtualMachine.Config .Memory VirtualMachine.Config .Settings VirtualMachine.Interac t.PowerOff VirtualMachine.Interac t.PowerOn VirtualMachine.Invent ory.CreateFromExisti ng VirtualMachine.Invent ory.Delete VirtualMachine.Provisi oning.Clone VirtualMachine.Provisi oning.DeployTemplat e
vSphere vCenter Datacenter	If the installation program creates the virtual machine folder. For user-provisioned infrastructure, VirtualMachine.Inventory.Create and VirtualMachine.Inventory.Delete privileges are optional if your cluster does not use the Machine API.	Resource.AssignVMT oPool VirtualMachine.Interac t.PowerOff VirtualMachine.Interac t.PowerOn VirtualMachine.Provisi oning.DeployTemplat e

Using OpenShift Container Platform with vMotion

If you intend on using vMotion in your vSphere environment, consider the following before installing an OpenShift Container Platform cluster.

- OpenShift Container Platform generally supports compute-only vMotion, where *generally* implies that you meet all VMware best practices for vMotion. To help ensure the uptime of your compute and control plane nodes, ensure that you follow the VMware best practices for vMotion, and use VMware anti-affinity rules to improve the availability of OpenShift Container Platform during maintenance or hardware issues.

For more information about vMotion and anti-affinity rules, see the VMware vSphere documentation for [vMotion networking requirements](#) and [VM anti-affinity rules](#).

- Using Storage vMotion can cause issues and is not supported. If you are using vSphere volumes in your pods, migrating a VM across datastores, either manually or through Storage vMotion, causes invalid references within OpenShift Container Platform persistent volume (PV) objects that can result in data loss.
- OpenShift Container Platform does not support selective migration of VMDKs across datastores, using datastore clusters for VM provisioning or for dynamic or static provisioning of PVs, or using a datastore that is part of a datastore cluster for dynamic or static provisioning of PVs.



IMPORTANT

You can specify the path of any datastore that exists in a datastore cluster. By default, Storage Distributed Resource Scheduler (SDRS), which uses Storage vMotion, is automatically enabled for a datastore cluster. Red Hat does not support Storage vMotion, so you must disable Storage DRS to avoid data loss issues for your OpenShift Container Platform cluster.

If you must specify VMs across multiple datastores, use a **datastore** object to specify a failure domain in your cluster's **install-config.yaml** configuration file. For more information, see "VMware vSphere region and zone enablement".

Cluster resources

When you deploy an OpenShift Container Platform cluster that uses installer-provisioned infrastructure, the installation program must be able to create several resources in your vCenter instance.

A standard OpenShift Container Platform installation creates the following vCenter resources:

- 1 Folder
- 1 Tag category
- 1 Tag
- Virtual machines:
 - 1 template
 - 1 temporary bootstrap node
 - 3 control plane nodes
 - 3 compute machines

Although these resources use 856 GB of storage, the bootstrap node is destroyed during the cluster installation process. A minimum of 800 GB of storage is required to use a standard cluster.

If you deploy more compute machines, the OpenShift Container Platform cluster will use more storage.

Cluster limits

Available resources vary between clusters. The number of possible clusters within a vCenter is limited primarily by available storage space and any limitations on the number of required resources. Be sure to consider both limitations to the vCenter resources that the cluster creates and the resources that you require to deploy a cluster, such as IP addresses and networks.

Networking requirements

Use Dynamic Host Configuration Protocol (DHCP) for the network and ensure that the DHCP server is configured to provide persistent IP addresses to the cluster machines.



NOTE


You do not need to use the DHCP for the network if you want to provision nodes with static IP addresses.

Configure the default gateway to use the DHCP server. All nodes must be in the same VLAN. You cannot scale the cluster using a second VLAN as a Day 2 operation.

You must use the Dynamic Host Configuration Protocol (DHCP) for the network and ensure that the DHCP server is configured to provide persistent IP addresses to the cluster machines. In the DHCP lease, you must configure the DHCP to use the default gateway. All nodes must be in the same VLAN. You cannot scale the cluster using a second VLAN as a Day 2 operation.

If you are installing to a restricted environment, the VM in your restricted network must have access to vCenter so that it can provision and manage nodes, persistent volume claims (PVCs), and other resources.

Additionally, you must create the following networking resources before you install the OpenShift Container Platform cluster:



NOTE

It is recommended that each OpenShift Container Platform node in the cluster must have access to a Network Time Protocol (NTP) server that is discoverable via DHCP. Installation is possible without an NTP server. However, asynchronous server clocks will cause errors, which NTP server prevents.

Required IP Addresses

For a network that uses DHCP, an installer-provisioned vSphere installation requires two static IP addresses:

- The **API** address is used to access the cluster API.
- The **Ingress** address is used for cluster ingress traffic.

You must provide these IP addresses to the installation program when you install the OpenShift Container Platform cluster.

DNS records

You must create DNS records for two static IP addresses in the appropriate DNS server for the vCenter instance that hosts your OpenShift Container Platform cluster. In each record, **<cluster_name>** is the cluster name and **<base_domain>** is the cluster base domain that you specify when you install the cluster. A complete DNS record takes the form: **<component>.<cluster_name>.<base_domain>..**

Table 23.6. Required DNS records

Component	Record	Description
-----------	--------	-------------

Component	Record	Description
API VIP	api.<cluster_name>.<base_domain>.	This DNS A/AAAA or CNAME (Canonical Name) record must point to the load balancer for the control plane machines. This record must be resolvable by both clients external to the cluster and from all the nodes within the cluster.
Ingress VIP	*.apps.<cluster_name>.<base_domain>.	A wildcard DNS A/AAAA or CNAME record that points to the load balancer that targets the machines that run the Ingress router pods, which are the worker nodes by default. This record must be resolvable by both clients external to the cluster and from all the nodes within the cluster.

Static IP addresses for vSphere nodes

You can provision bootstrap, control plane, and compute nodes to be configured with static IP addresses in environments where Dynamic Host Configuration Protocol (DHCP) does not exist. To configure this environment, you must provide values to the **platform.vsphere.hosts.role** parameter in the **install-config.yaml** file.

By default, the installation program is configured to use the DHCP for the network, but this network has limited configurable capabilities.

After you define one or more machine pools in your **install-config.yaml** file, you can define network definitions for nodes on your network. Ensure that the number of network definitions matches the number of machine pools that you configured for your cluster.

The following example shows a network configuration for a node with the role **compute**:

```

---
platform:
  vsphere:
    hosts:
      - role: compute 1
      networkDevice:
        ipAddrs:
          - 192.168.204.10/24 2
        gateway: 192.168.204.1 3
        nameservers:
          - 192.168.204.1 4
---

```

- 1 Valid network definition values include **bootstrap**, **control-plane**, and **compute**. You must list at least one **bootstrap** network definition in your **install-config.yaml** configuration file.
- 2 Lists IPv4, IPv6, or both IP addresses that the installation program passes to the network interface. The machine API controller assigns all configured IP addresses to the default network interface.
- 3 The default gateway for the network interface.
- 4 Lists up to 3 DNS nameservers.

After you deployed your cluster to run nodes with static IP addresses, you can scale a machine to use one of these static IP addresses. Additionally, you can use a machine set to configure a machine to use one of the configured static IP addresses.


Additional resources

- [Scaling machines to use static IP addresses](#)
- [Using a machine set to scale machines with configured static IP addresses](#)

23.2.2. Preparing to install a cluster using installer-provisioned infrastructure

You prepare to install an OpenShift Container Platform cluster on vSphere by completing the following steps:

- Downloading the installation program.



NOTE

If you are installing in a disconnected environment, you extract the installation program from the mirrored content. For more information, see [Mirroring images for a disconnected installation](#).

- Installing the OpenShift CLI (**oc**).



NOTE

If you are installing in a disconnected environment, install **oc** to the mirror host.


- Generating an SSH key pair. You can use this key pair to authenticate into the OpenShift Container Platform cluster's nodes after it is deployed.
- Adding your vCenter's trusted root CA certificates to your system trust.

23.2.2.1. Obtaining the installation program

Before you install OpenShift Container Platform, download the installation file on the host you are using for installation.

Prerequisites

- You have a machine that runs Linux, for example Red Hat Enterprise Linux 8, with 500 MB of local disk space.




IMPORTANT

If you attempt to run the installation program on macOS, a known issue related to the **golang** compiler causes the installation of the OpenShift Container Platform cluster to fail. For more information about this issue, see the section named "Known Issues" in the *OpenShift Container Platform 4.16 release notes* document.


Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select your infrastructure provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.



IMPORTANT

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.



IMPORTANT

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.


4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#). This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

23.2.2.2. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.



IMPORTANT

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**.
To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.



NOTE

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your PATH.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

23.2.2.3. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.



NOTE


You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.



NOTE

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.


2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.



NOTE

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```



NOTE

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1 Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

23.2.2.4. Adding vCenter root CA certificates to your system trust

Because the installation program requires access to your vCenter's API, you must add your vCenter's trusted root CA certificates to your system trust before you install an OpenShift Container Platform cluster.

Procedure

1. From the vCenter home page, download the vCenter's root CA certificates. Click **Download trusted root CA certificates** in the vSphere Web Services SDK section. The `<vCenter>/certs/download.zip` file downloads.
2. Extract the compressed file that contains the vCenter root CA certificates. The contents of the compressed file resemble the following file structure:

```
certs
├── lin
│   ├── 108f4d17.0
│   ├── 108f4d17.r1
│   ├── 7e757f6a.0
│   ├── 8e4f8471.0
│   └── 8e4f8471.r0
├── mac
│   ├── 108f4d17.0
│   ├── 108f4d17.r1
│   ├── 7e757f6a.0
│   ├── 8e4f8471.0
│   └── 8e4f8471.r0
└── win
    ├── 108f4d17.0.crt
    ├── 108f4d17.r1.crl
    ├── 7e757f6a.0.crt
    ├── 8e4f8471.0.crt
    └── 8e4f8471.r0.crl
```

3 directories, 15 files

3. Add the files for your operating system to the system trust. For example, on a Fedora operating system, run the following command:

```
# cp certs/lin/* /etc/pki/ca-trust/source/anchors
```

- Update your system trust. For example, on a Fedora operating system, run the following command:

```
# update-ca-trust extract
```

23.2.3. Installing a cluster on vSphere

In OpenShift Container Platform version 4.16, you can install a cluster on your VMware vSphere instance by using installer-provisioned infrastructure.




NOTE

OpenShift Container Platform supports deploying a cluster to a single VMware vCenter only. Deploying a cluster with machines/machine sets on multiple vCenters is not supported.

23.2.3.1. Prerequisites

- You have completed the tasks in [Preparing to install a cluster using installer-provisioned infrastructure](#).
- You reviewed your VMware platform licenses. Red Hat does not place any restrictions on your VMware licenses, but some VMware infrastructure components require licensing.
- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You provisioned [persistent storage](#) for your cluster. To deploy a private image registry, your storage must provide **ReadWriteMany** access modes.
- The OpenShift Container Platform installer requires access to port 443 on the vCenter and ESXi hosts. You verified that port 443 is accessible.
- If you use a firewall, you confirmed with the administrator that port 443 is accessible. Control plane nodes must be able to reach vCenter and ESXi hosts on port 443 for the installation to succeed.
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.



NOTE

Be sure to also review this site list if you are configuring a proxy.

23.2.3.2. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.




IMPORTANT

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

23.2.3.3. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.




IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.
- Optional: Before you create the cluster, configure an external load balancer in place of the default load balancer.



IMPORTANT

You do not need to specify API and Ingress static addresses for your installation program. If you choose this configuration, you must take additional actions to define network targets that accept an IP address from each referenced vSphere subnet. See the section "Configuring a user-managed load balancer".

Procedure

1. Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1  
--log-level=info 2
```

- 1 For **<installation_directory>**, specify the directory name to store the files that the installation program creates.
- 2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

When specifying the directory:

- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
- Use an empty directory. Some installation assets, such as bootstrap X.509 certificates, have short expiration intervals, therefore you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

2. Provide values at the prompts:


- a. Optional: Select an SSH key to use to access your cluster machines.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- b. Select **vsphere** as the platform to target.
- c. Specify the name of your vCenter instance.
- d. Specify the user name and password for the vCenter account that has the required permissions to create the cluster.
The installation program connects to your vCenter instance.



IMPORTANT

Some VMware vCenter Single Sign-On (SSO) environments with Active Directory (AD) integration might primarily require you to use the traditional login method, which requires the **<domain>**\ construct.

To ensure that vCenter account permission checks complete properly, consider using the User Principal Name (UPN) login method, such as **<username>@<fully_qualified_domainname>**.

- e. Select the data center in your vCenter instance to connect to.
- f. Select the default vCenter datastore to use.



NOTE

Datastore and cluster names cannot exceed 60 characters; therefore, ensure the combined string length does not exceed the 60 character limit.

- g. Select the vCenter cluster to install the OpenShift Container Platform cluster in. The installation program uses the root resource pool of the vSphere cluster as the default resource pool.
- h. Select the network in the vCenter instance that contains the virtual IP addresses and DNS records that you configured.
- i. Enter the virtual IP address that you configured for control plane API access.
- j. Enter the virtual IP address that you configured for cluster ingress.
- k. Enter the base domain. This base domain must be the same one that you used in the DNS records that you configured.
- l. Enter a descriptive name for your cluster. The cluster name must be the same one that you used in the DNS records that you configured.



NOTE


Datastore and cluster names cannot exceed 60 characters; therefore, ensure the combined string length does not exceed the 60 character limit.

- m. Paste the [pull secret from Red Hat OpenShift Cluster Manager](#) .

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to `<installation_directory>/openshift_install.log`.



IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```




IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

23.2.3.4. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

23.2.3.5. Creating registry storage

After you install the cluster, you must create storage for the registry Operator.

23.2.3.5.1. Image registry removed during installation

On platforms that do not provide shareable object storage, the OpenShift Image Registry Operator bootstraps itself as **Removed**. This allows **openshift-installer** to complete installations on these platform types.

After installation, you must edit the Image Registry Operator configuration to switch the **managementState** from **Removed** to **Managed**. When this has completed, you must configure storage.

23.2.3.5.2. Image registry storage configuration

The Image Registry Operator is not initially available for platforms that do not provide default storage. After installation, you must configure your registry to use storage so that the Registry Operator is made available.

Instructions are shown for configuring a persistent volume, which is required for production clusters. Where applicable, instructions are shown for configuring an empty directory as the storage location, which is available for only non-production clusters.


Additional instructions are provided for allowing the image registry to use block storage types by using the **Recreate** rollout strategy during upgrades.

23.2.3.5.2.1. Configuring registry storage for VMware vSphere

As a cluster administrator, following installation you must configure your registry to use storage.

Prerequisites

- Cluster administrator permissions.
- A cluster on VMware vSphere.
- Persistent storage provisioned for your cluster, such as Red Hat OpenShift Data Foundation.



IMPORTANT

OpenShift Container Platform supports **ReadWriteOnce** access for image registry storage when you have only one replica. **ReadWriteOnce** access also requires that the registry uses the **Recreate** rollout strategy. To deploy an image registry that supports high availability with two or more replicas, **ReadWriteMany** access is required.

- Must have "100Gi" capacity.




IMPORTANT

Testing shows issues with using the NFS server on RHEL as storage backend for core services. This includes the OpenShift Container Registry and Quay, Prometheus for monitoring storage, and Elasticsearch for logging storage. Therefore, using RHEL NFS to back PVs used by core services is not recommended.

Other NFS implementations on the marketplace might not have these issues. Contact the individual NFS implementation vendor for more information on any testing that was possibly completed against these OpenShift Container Platform core components.

Procedure

1. To configure your registry to use storage, change the **spec.storage.pvc** in the **configs.imageregistry/cluster** resource.

**NOTE**


When you use shared storage, review your security settings to prevent outside access.

2. Verify that you do not have a registry pod:

```
$ oc get pod -n openshift-image-registry -l docker-registry=default
```

Example output

```
No resources found in openshift-image-registry namespace
```

**NOTE**

If you do have a registry pod in your output, you do not need to continue with this procedure.

3. Check the registry configuration:

```
$ oc edit configs.imageregistry.operator.openshift.io
```

Example output

```
storage:
  pvc:
    claim: 1
```

- 1** Leave the **claim** field blank to allow the automatic creation of an **image-registry-storage** persistent volume claim (PVC). The PVC is generated based on the default storage class. However, be aware that the default storage class might provide ReadWriteOnce (RWO) volumes, such as a RADOS Block Device (RBD), which can cause issues when you replicate to more than one replica.

4. Check the **clusteroperator** status:

```
$ oc get clusteroperator image-registry
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED
image-registry	4.7	True	False	False

23.2.3.5.2.2. Configuring block registry storage for VMware vSphere

To allow the image registry to use block storage types such as vSphere Virtual Machine Disk (VMDK) during upgrades as a cluster administrator, you can use the **Recreate** rollout strategy.



IMPORTANT

Block storage volumes are supported but not recommended for use with image registry on production clusters. An installation where the registry is configured on block storage is not highly available because the registry cannot have more than one replica.

Procedure

1. Enter the following command to set the image registry storage as a block storage type, patch the registry so that it uses the **Recreate** rollout strategy, and runs with only **1** replica:

```
$ oc patch config.imageregistry.operator.openshift.io/cluster --type=merge -p '{"spec": {"rolloutStrategy": "Recreate", "replicas": 1}}'
```

2. Provision the PV for the block storage device, and create a PVC for that volume. The requested block volume uses the ReadWriteOnce (RWO) access mode.

- a. Create a **pvc.yaml** file with the following contents to define a VMware vSphere **PersistentVolumeClaim** object:

```
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: image-registry-storage 1
  namespace: openshift-image-registry 2
spec:
  accessModes:
  - ReadWriteOnce 3
  resources:
    requests:
      storage: 100Gi 4
```

- 1 A unique name that represents the **PersistentVolumeClaim** object.
- 2 The namespace for the **PersistentVolumeClaim** object, which is **openshift-image-registry**.
- 3 The access mode of the persistent volume claim. With **ReadWriteOnce**, the volume can be mounted with read and write permissions by a single node.
- 4 The size of the persistent volume claim.

- b. Enter the following command to create the **PersistentVolumeClaim** object from the file:

```
$ oc create -f pvc.yaml -n openshift-image-registry
```

3. Enter the following command to edit the registry configuration so that it references the correct PVC:

```
$ oc edit config.imageregistry.operator.openshift.io -o yaml
```

Example output

```
storage:
  pvc:
    claim: 1
```

- 1 By creating a custom PVC, you can leave the **claim** field blank for the default automatic creation of an **image-registry-storage** PVC.

For instructions about configuring registry storage so that it references the correct PVC, see [Configuring the registry for vSphere](#).

23.2.3.6. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service

23.2.3.7. Next steps

- [Customize your cluster](#).
- If necessary, you can [opt out of remote health reporting](#).
- [Set up your registry and configure registry storage](#).
- Optional: [View the events from the vSphere Problem Detector Operator](#) to determine if the cluster has permission or storage configuration issues.

23.2.4. Installing a cluster on vSphere with customizations

In OpenShift Container Platform version 4.16, you can install a cluster on your VMware vSphere instance by using installer-provisioned infrastructure. To customize the installation, you modify parameters in the **install-config.yaml** file before you install the cluster.




NOTE

OpenShift Container Platform supports deploying a cluster to a single VMware vCenter only. Deploying a cluster with machines/machine sets on multiple vCenters is not supported.

23.2.4.1. Prerequisites

- You have completed the tasks in [Preparing to install a cluster using installer-provisioned infrastructure](#).

- You reviewed your VMware platform licenses. Red Hat does not place any restrictions on your VMware licenses, but some VMware infrastructure components require licensing.
- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You provisioned [persistent storage](#) for your cluster. To deploy a private image registry, your storage must provide **ReadWriteMany** access modes.
- The OpenShift Container Platform installer requires access to port 443 on the vCenter and ESXi hosts. You verified that port 443 is accessible.
- If you use a firewall, you confirmed with the administrator that port 443 is accessible. Control plane nodes must be able to reach vCenter and ESXi hosts on port 443 for the installation to succeed.
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.



NOTE


Be sure to also review this site list if you are configuring a proxy.

23.2.4.2. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.



IMPORTANT

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

23.2.4.3. VMware vSphere region and zone enablement

You can deploy an OpenShift Container Platform cluster to multiple vSphere datacenters that run in a single VMware vCenter. Each datacenter can run multiple clusters. This configuration reduces the risk of a hardware failure or network outage that can cause your cluster to fail. To enable regions and zones, you must define multiple failure domains for your OpenShift Container Platform cluster.



IMPORTANT

The VMware vSphere region and zone enablement feature requires the vSphere Container Storage Interface (CSI) driver as the default storage driver in the cluster. As a result, the feature is only available on a newly installed cluster.


For a cluster that was upgraded from a previous release, you must enable CSI automatic migration for the cluster. You can then configure multiple regions and zones for the upgraded cluster.

The default installation configuration deploys a cluster to a single vSphere datacenter. If you want to deploy a cluster to multiple vSphere datacenters, you must create an installation configuration file that enables the region and zone feature.

The default **install-config.yaml** file includes **vccenters** and **failureDomains** fields, where you can specify multiple vSphere datacenters and clusters for your OpenShift Container Platform cluster. You can leave these fields blank if you want to install an OpenShift Container Platform cluster in a vSphere environment that consists of single datacenter.

The following list describes terms associated with defining zones and regions for your cluster:

- **Failure domain:** Establishes the relationships between a region and zone. You define a failure domain by using vCenter objects, such as a **datastore** object. A failure domain defines the vCenter location for OpenShift Container Platform cluster nodes.
- **Region:** Specifies a vCenter datacenter. You define a region by using a tag from the **openshift-region** tag category.
- **Zone:** Specifies a vCenter cluster. You define a zone by using a tag from the **openshift-zone** tag category.



NOTE

If you plan on specifying more than one failure domain in your **install-config.yaml** file, you must create tag categories, zone tags, and region tags in advance of creating the configuration file.

You must create a vCenter tag for each vCenter datacenter, which represents a region. Additionally, you must create a vCenter tag for each cluster that runs in a datacenter, which represents a zone. After you create the tags, you must attach each tag to their respective datacenters and clusters.

The following table outlines an example of the relationship among regions, zones, and tags for a configuration with multiple vSphere datacenters running in a single VMware vCenter.

Datacenter (region)	Cluster (zone)	Tags
us-east	us-east-1	us-east-1a
		us-east-1b
	us-east-2	us-east-2a
		us-east-2b

Datacenter (region)	Cluster (zone)	Tags
us-west	us-west-1	us-west-1a
		us-west-1b
	us-west-2	us-west-2a
		us-west-2b

Additional resources

- [Additional VMware vSphere configuration parameters](#)
- [Deprecated VMware vSphere configuration parameters](#)
- [vSphere automatic migration](#)
- [VMware vSphere CSI Driver Operator](#)

23.2.4.4. Creating the installation configuration file

You can customize the OpenShift Container Platform cluster you install on VMware vSphere.

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create the **install-config.yaml** file.
 - a. Change to the directory that contains the installation program and run the following command:

```

$ ./openshift-install create install-config --dir <installation_directory> 1

```

- 1** For **<installation_directory>**, specify the directory name to store the files that the installation program creates.

When specifying the directory:

- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
- Use an empty directory. Some installation assets, such as bootstrap X.509 certificates, have short expiration intervals, therefore you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

- b. At the prompts, provide the configuration details for your cloud:
- i. Optional: Select an SSH key to use to access your cluster machines.

**NOTE**

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- ii. Select **vsphere** as the platform to target.
- iii. Specify the name of your vCenter instance.
- iv. Specify the user name and password for the vCenter account that has the required permissions to create the cluster.
The installation program connects to your vCenter instance.
- v. Select the data center in your vCenter instance to connect to.

**NOTE**

After you create the installation configuration file, you can modify the file to create a multiple vSphere datacenters environment. This means that you can deploy an OpenShift Container Platform cluster to multiple vSphere datacenters that run in a single VMware vCenter. For more information about creating this environment, see the section named *VMware vSphere region and zone enablement*.

- vi. Select the default vCenter datastore to use.

**WARNING**

You can specify the path of any datastore that exists in a datastore cluster. By default, Storage Distributed Resource Scheduler (SDRS), which uses Storage vMotion, is automatically enabled for a datastore cluster. Red Hat does not support Storage vMotion, so you must disable Storage DRS to avoid data loss issues for your OpenShift Container Platform cluster.

You cannot specify more than one datastore path. If you must specify VMs across multiple datastores, use a **datastore** object to specify a failure domain in your cluster's **install-config.yaml** configuration file. For more information, see "VMware vSphere region and zone enablement".


- vii. Select the vCenter cluster to install the OpenShift Container Platform cluster in. The installation program uses the root resource pool of the vSphere cluster as the default resource pool.

- viii. Select the network in the vCenter instance that contains the virtual IP addresses and DNS records that you configured.
 - ix. Enter the virtual IP address that you configured for control plane API access.
 - x. Enter the virtual IP address that you configured for cluster ingress.
 - xi. Enter the base domain. This base domain must be the same one that you used in the DNS records that you configured.
 - xii. Enter a descriptive name for your cluster.
The cluster name you enter must match the cluster name you specified when configuring the DNS records.
2. Modify the **install-config.yaml** file. You can find more information about the available parameters in the "Installation configuration parameters" section.

**NOTE**

If you are installing a three-node cluster, be sure to set the **compute.replicas** parameter to **0**. This ensures that the cluster's control planes are schedulable. For more information, see "Installing a three-node cluster on vSphere".

3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.

**IMPORTANT**

The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

Additional resources

- [Installation configuration parameters](#)

23.2.4.4.1. Sample install-config.yaml file for an installer-provisioned VMware vSphere cluster

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.

```

apiVersion: v1
baseDomain: example.com 1
compute: 2
- architecture: amd64
  name: <worker_node>
  platform: {}
  replicas: 3
controlPlane: 3
  architecture: amd64
  name: <parent_node>
  platform: {}
  replicas: 3
metadata:
  creationTimestamp: null
  name: test 4

```

```

platform:
vsphere: 5
  apiVIPs:
    - 10.0.0.1
  failureDomains: 6
    - name: <failure_domain_name>
      region: <default_region_name>
      server: <fully_qualified_domain_name>
  topology:
    computeCluster: "/<datacenter>/host/<cluster>"
    datacenter: <datacenter>
    datastore: "/<datacenter>/datastore/<datastore>" 7
    networks:
      - <VM_Network_name>
    resourcePool: "/<datacenter>/host/<cluster>/Resources/<resourcePool>" 8
    folder: "/<datacenter_name>/vm/<folder_name>/<subfolder_name>"
    tagIDs: 9
      - <tag_id> 10
    zone: <default_zone_name>
  ingressVIPs:
    - 10.0.0.2
  vcenters:
    - datacenters:
      - <datacenter>
      password: <password>
      port: 443
      server: <fully_qualified_domain_name>
      user: administrator@vsphere.local
    diskType: thin 11
  fips: false
  pullSecret: '{"auths": ...}'
  sshKey: 'ssh-ed25519 AAAA...'

```

- 1 The base domain of the cluster. All DNS records must be sub-domains of this base and include the cluster name.
- 2 3 The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.
- 4 The cluster name that you specified in your DNS records.
- 5 Optional: Provides additional configuration for the machine pool parameters for the compute and control plane machines.
- 6 Establishes the relationships between a region and zone. You define a failure domain by using vCenter objects, such as a **datastore** object. A failure domain defines the vCenter location for OpenShift Container Platform cluster nodes.
- 7 The path to the vSphere datastore that holds virtual machine files, templates, and ISO images.



IMPORTANT

You can specify the path of any datastore that exists in a datastore cluster. By default, Storage vMotion is automatically enabled for a datastore cluster. Red Hat does not support Storage vMotion, so you must disable Storage vMotion to avoid data loss issues for your OpenShift Container Platform cluster.

If you must specify VMs across multiple datastores, use a **datastore** object to specify a failure domain in your cluster's **install-config.yaml** configuration file. For more information, see "VMware vSphere region and zone enablement".

- 8 Optional: Provides an existing resource pool for machine creation. If you do not specify a value, the installation program uses the root resource pool of the vSphere cluster.
- 9 Optional: Each VM created by OpenShift Container Platform is assigned a unique tag that is specific to the cluster. The assigned tag enables the installation program to identify and remove the associated VMs when a cluster is decommissioned. You can list up to ten additional tag IDs to be attached to the VMs provisioned by the installation program.
- 10 The ID of the tag to be associated by the installation program. For example, **urn:vmomi:InventoryServiceTag:208e713c-cae3-4b7f-918e-4051ca7d1f97:GLOBAL**. For more information about determining the tag ID, see the [vSphere Tags and Attributes documentation](#).
- 11 The vSphere disk provisioning method.



NOTE

In OpenShift Container Platform 4.12 and later, the **apiVIP** and **ingressVIP** configuration settings are deprecated. Instead, use a list format to enter values in the **apiVIPs** and **ingressVIPs** configuration settings.

23.2.4.4.2. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations. You must include vCenter's IP address and the IP range that you use for its machines.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.

**NOTE**

The installation program does not support the proxy **readinessEndpoints** field.

**NOTE**

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$. /openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.

**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

23.2.4.4.3. Configuring regions and zones for a VMware vCenter

You can modify the default installation configuration file, so that you can deploy an OpenShift Container Platform cluster to multiple vSphere datacenters that run in a single VMware vCenter.


The default **install-config.yaml** file configuration from the previous release of OpenShift Container Platform is deprecated. You can continue to use the deprecated default configuration, but the **openshift-installer** will prompt you with a warning message that indicates the use of deprecated fields in the configuration file.

**IMPORTANT**

The example uses the **govc** command. The **govc** command is an open source command available from VMware; it is not available from Red Hat. The Red Hat support team does not maintain the **govc** command. Instructions for downloading and installing **govc** are found on the VMware documentation website

Prerequisites


- You have an existing **install-config.yaml** installation configuration file.

**IMPORTANT**

You must specify at least one failure domain for your OpenShift Container Platform cluster, so that you can provision datacenter objects for your VMware vCenter server. Consider specifying multiple failure domains if you need to provision virtual machine nodes in different datacenters, clusters, datastores, and other components. To enable regions and zones, you must define multiple failure domains for your OpenShift Container Platform cluster.

Procedure

1. Enter the following **govc** command-line tool commands to create the **openshift-region** and **openshift-zone** vCenter tag categories:



IMPORTANT

If you specify different names for the **openshift-region** and **openshift-zone** vCenter tag categories, the installation of the OpenShift Container Platform cluster fails.

```
$ govc tags.category.create -d "OpenShift region" openshift-region
```

```
$ govc tags.category.create -d "OpenShift zone" openshift-zone
```

2. To create a region tag for each region vSphere datacenter where you want to deploy your cluster, enter the following command in your terminal:

```
$ govc tags.create -c <region_tag_category> <region_tag>
```

3. To create a zone tag for each vSphere cluster where you want to deploy your cluster, enter the following command:

```
$ govc tags.create -c <zone_tag_category> <zone_tag>
```

4. Attach region tags to each vCenter datacenter object by entering the following command:

```
$ govc tags.attach -c <region_tag_category> <region_tag_1> /<datacenter_1>
```

5. Attach the zone tags to each vCenter datacenter object by entering the following command:

```
$ govc tags.attach -c <zone_tag_category> <zone_tag_1> /<datacenter_1>/host/vcs-mdcnc-workload-1
```

6. Change to the directory that contains the installation program and initialize the cluster deployment according to your chosen installation requirements.

Sample install-config.yaml file with multiple datacenters defined in a vSphere center

```
---
compute:
---
vsphere:
  zones:
    - "<machine_pool_zone_1>"
    - "<machine_pool_zone_2>"
---
controlPlane:
---
vsphere:
  zones:
    - "<machine_pool_zone_1>"
    - "<machine_pool_zone_2>"
---
```

```

platform:
  vsphere:
    vcenters:
    ---
  datacenters:
    - <datacenter1_name>
    - <datacenter2_name>
  failureDomains:
    - name: <machine_pool_zone_1>
      region: <region_tag_1>
      zone: <zone_tag_1>
      server: <fully_qualified_domain_name>
    topology:
      datacenter: <datacenter1>
      computeCluster: "/<datacenter1>/host/<cluster1>"
      networks:
      - <VM_Network1_name>
      datastore: "/<datacenter1>/datastore/<datastore1>"
      resourcePool: "/<datacenter1>/host/<cluster1>/Resources/<resourcePool1>"
      folder: "/<datacenter1>/vm/<folder1>"
    - name: <machine_pool_zone_2>
      region: <region_tag_2>
      zone: <zone_tag_2>
      server: <fully_qualified_domain_name>
    topology:
      datacenter: <datacenter2>
      computeCluster: "/<datacenter2>/host/<cluster2>"
      networks:
      - <VM_Network2_name>
      datastore: "/<datacenter2>/datastore/<datastore2>"
      resourcePool: "/<datacenter2>/host/<cluster2>/Resources/<resourcePool2>"
      folder: "/<datacenter2>/vm/<folder2>"
    ---

```

23.2.4.5. Services for a user-managed load balancer

You can configure an OpenShift Container Platform cluster to use a user-managed load balancer in place of the default load balancer.



IMPORTANT

Configuring a user-managed load balancer depends on your vendor's load balancer.


The information and examples in this section are for guideline purposes only. Consult the vendor documentation for more specific information about the vendor's load balancer.

Red Hat supports the following services for a user-managed load balancer:

- Ingress Controller
- OpenShift API
- OpenShift MachineConfig API

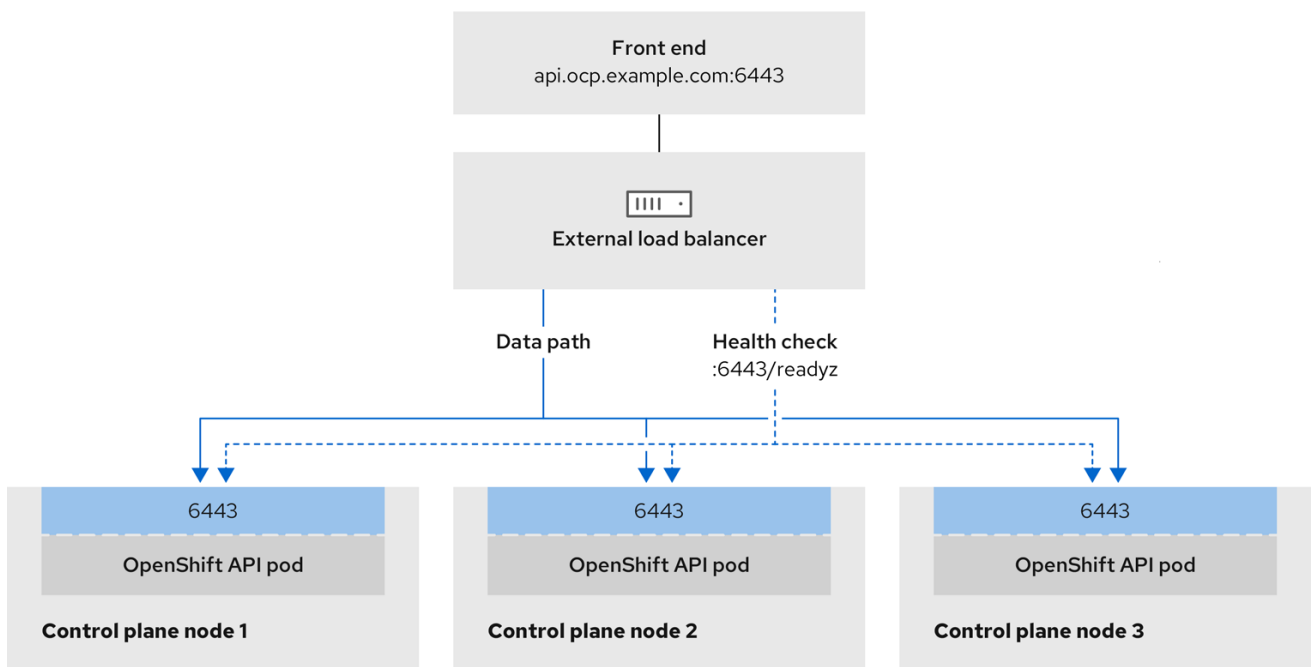
You can choose whether you want to configure one or all of these services for a user-managed load balancer. Configuring only the Ingress Controller service is a common configuration option. To better understand each service, view the following diagrams:

Figure 23.1. Example network workflow that shows an Ingress Controller operating in an OpenShift Container Platform environment



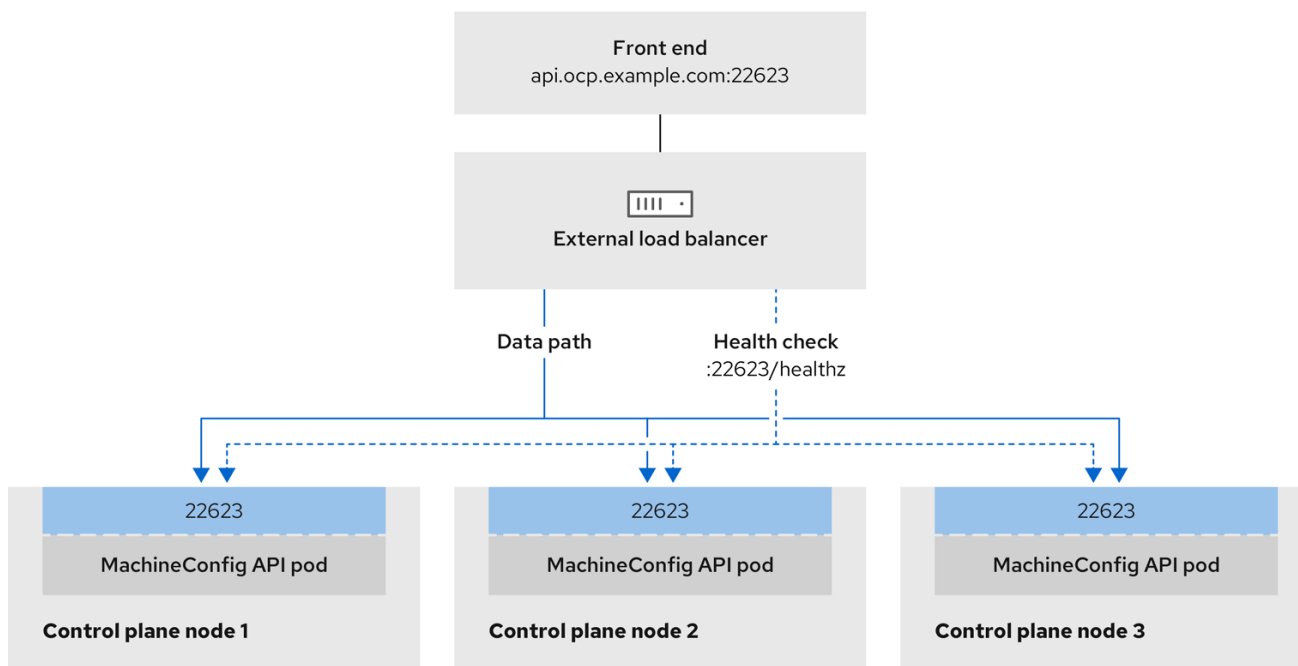
496_OpenShift_1223

Figure 23.2. Example network workflow that shows an OpenShift API operating in an OpenShift Container Platform environment



496_OpenShift_1223

Figure 23.3. Example network workflow that shows an OpenShift MachineConfig API operating in an OpenShift Container Platform environment



496_OpenShift_I223

The following configuration options are supported for user-managed load balancers:

- Use a node selector to map the Ingress Controller to a specific set of nodes. You must assign a static IP address to each node in this set, or configure each node to receive the same IP address from the Dynamic Host Configuration Protocol (DHCP). Infrastructure nodes commonly receive this type of configuration.
- Target all IP addresses on a subnet. This configuration can reduce maintenance overhead, because you can create and destroy nodes within those networks without reconfiguring the load balancer targets. If you deploy your ingress pods by using a machine set on a smaller network, such as a `/27` or `/28`, you can simplify your load balancer targets.

TIP

You can list all IP addresses that exist in a network by checking the machine config pool's resources.


Before you configure a user-managed load balancer for your OpenShift Container Platform cluster, consider the following information:

- For a front-end IP address, you can use the same IP address for the front-end IP address, the Ingress Controller's load balancer, and API load balancer. Check the vendor's documentation for this capability.
- For a back-end IP address, ensure that an IP address for an OpenShift Container Platform control plane node does not change during the lifetime of the user-managed load balancer. You can achieve this by completing one of the following actions:
 - Assign a static IP address to each control plane node.

- Configure each node to receive the same IP address from the DHCP every time the node requests a DHCP lease. Depending on the vendor, the DHCP lease might be in the form of an IP reservation or a static DHCP assignment.
- Manually define each node that runs the Ingress Controller in the user-managed load balancer for the Ingress Controller back-end service. For example, if the Ingress Controller moves to an undefined node, a connection outage can occur.

23.2.4.5.1. Configuring a user-managed load balancer

You can configure an OpenShift Container Platform cluster to use a user-managed load balancer in place of the default load balancer.



IMPORTANT

Before you configure a user-managed load balancer, ensure that you read the "Services for a user-managed load balancer" section.

Read the following prerequisites that apply to the service that you want to configure for your user-managed load balancer.



NOTE

MetalLB, which runs on a cluster, functions as a user-managed load balancer.

OpenShift API prerequisites

- You defined a front-end IP address.
- TCP ports 6443 and 22623 are exposed on the front-end IP address of your load balancer. Check the following items:
 - Port 6443 provides access to the OpenShift API service.
 - Port 22623 can provide ignition startup configurations to nodes.
- The front-end IP address and port 6443 are reachable by all users of your system with a location external to your OpenShift Container Platform cluster.
- The front-end IP address and port 22623 are reachable only by OpenShift Container Platform nodes.
- The load balancer backend can communicate with OpenShift Container Platform control plane nodes on port 6443 and 22623.

Ingress Controller prerequisites

- You defined a front-end IP address.
- TCP ports 443 and 80 are exposed on the front-end IP address of your load balancer.
- The front-end IP address, port 80 and port 443 are be reachable by all users of your system with a location external to your OpenShift Container Platform cluster.

- The front-end IP address, port 80 and port 443 are reachable to all nodes that operate in your OpenShift Container Platform cluster.
- The load balancer backend can communicate with OpenShift Container Platform nodes that run the Ingress Controller on ports 80, 443, and 1936.

Prerequisite for health check URL specifications

You can configure most load balancers by setting health check URLs that determine if a service is available or unavailable. OpenShift Container Platform provides these health checks for the OpenShift API, Machine Configuration API, and Ingress Controller backend services.

The following examples show health check specifications for the previously listed backend services:

Example of a Kubernetes API health check specification

```
Path: HTTPS:6443/readyz
Healthy threshold: 2
Unhealthy threshold: 2
Timeout: 10
Interval: 10
```

Example of a Machine Config API health check specification

```
Path: HTTPS:22623/healthz
Healthy threshold: 2
Unhealthy threshold: 2
Timeout: 10
Interval: 10
```

Example of an Ingress Controller health check specification

```
Path: HTTP:1936/healthz/ready
Healthy threshold: 2
Unhealthy threshold: 2
Timeout: 5
Interval: 10
```

Procedure

1. Configure the HAProxy Ingress Controller, so that you can enable access to the cluster from your load balancer on ports 6443, 22623, 443, and 80. Depending on your needs, you can specify the IP address of a single subnet or IP addresses from multiple subnets in your HAProxy configuration.

Example HAProxy configuration with one listed subnet

```
# ...
listen my-cluster-api-6443
    bind 192.168.1.100:6443
    mode tcp
    balance roundrobin
    option httpchk
    http-check connect
    http-check send meth GET uri /readyz
```

```

http-check expect status 200
  server my-cluster-master-2 192.168.1.101:6443 check inter 10s rise 2 fall 2
  server my-cluster-master-0 192.168.1.102:6443 check inter 10s rise 2 fall 2
  server my-cluster-master-1 192.168.1.103:6443 check inter 10s rise 2 fall 2

listen my-cluster-machine-config-api-22623
  bind 192.168.1.100:22623
  mode tcp
  balance roundrobin
  option httpchk
  http-check connect
  http-check send meth GET uri /healthz
  http-check expect status 200
    server my-cluster-master-2 192.168.1.101:22623 check inter 10s rise 2 fall 2
    server my-cluster-master-0 192.168.1.102:22623 check inter 10s rise 2 fall 2
    server my-cluster-master-1 192.168.1.103:22623 check inter 10s rise 2 fall 2

listen my-cluster-apps-443
  bind 192.168.1.100:443
  mode tcp
  balance roundrobin
  option httpchk
  http-check connect
  http-check send meth GET uri /healthz/ready
  http-check expect status 200
    server my-cluster-worker-0 192.168.1.111:443 check port 1936 inter 10s rise 2 fall 2
    server my-cluster-worker-1 192.168.1.112:443 check port 1936 inter 10s rise 2 fall 2
    server my-cluster-worker-2 192.168.1.113:443 check port 1936 inter 10s rise 2 fall 2

listen my-cluster-apps-80
  bind 192.168.1.100:80
  mode tcp
  balance roundrobin
  option httpchk
  http-check connect
  http-check send meth GET uri /healthz/ready
  http-check expect status 200
    server my-cluster-worker-0 192.168.1.111:80 check port 1936 inter 10s rise 2 fall 2
    server my-cluster-worker-1 192.168.1.112:80 check port 1936 inter 10s rise 2 fall 2
    server my-cluster-worker-2 192.168.1.113:80 check port 1936 inter 10s rise 2 fall 2
# ...

```

Example HAProxy configuration with multiple listed subnets

```

# ...
listen api-server-6443
  bind *:6443
  mode tcp
    server master-00 192.168.83.89:6443 check inter 1s
    server master-01 192.168.84.90:6443 check inter 1s
    server master-02 192.168.85.99:6443 check inter 1s
    server bootstrap 192.168.80.89:6443 check inter 1s

listen machine-config-server-22623
  bind *:22623
  mode tcp

```

```

server master-00 192.168.83.89:22623 check inter 1s
server master-01 192.168.84.90:22623 check inter 1s
server master-02 192.168.85.99:22623 check inter 1s
server bootstrap 192.168.80.89:22623 check inter 1s

listen ingress-router-80
  bind *:80
  mode tcp
  balance source
    server worker-00 192.168.83.100:80 check inter 1s
    server worker-01 192.168.83.101:80 check inter 1s

listen ingress-router-443
  bind *:443
  mode tcp
  balance source
    server worker-00 192.168.83.100:443 check inter 1s
    server worker-01 192.168.83.101:443 check inter 1s

listen ironic-api-6385
  bind *:6385
  mode tcp
  balance source
    server master-00 192.168.83.89:6385 check inter 1s
    server master-01 192.168.84.90:6385 check inter 1s
    server master-02 192.168.85.99:6385 check inter 1s
    server bootstrap 192.168.80.89:6385 check inter 1s

listen inspector-api-5050
  bind *:5050
  mode tcp
  balance source
    server master-00 192.168.83.89:5050 check inter 1s
    server master-01 192.168.84.90:5050 check inter 1s
    server master-02 192.168.85.99:5050 check inter 1s
    server bootstrap 192.168.80.89:5050 check inter 1s
# ...

```

2. Use the **curl** CLI command to verify that the user-managed load balancer and its resources are operational:
 - a. Verify that the cluster machine configuration API is accessible to the Kubernetes API server resource, by running the following command and observing the response:

```
$ curl https://<loadbalancer_ip_address>:6443/version --insecure
```

If the configuration is correct, you receive a JSON object in response:

```
{
  "major": "1",
  "minor": "11+",
  "gitVersion": "v1.11.0+ad103ed",
  "gitCommit": "ad103ed",
  "gitTreeState": "clean",
  "buildDate": "2019-01-09T06:44:10Z",
  "goVersion": "go1.10.3",
}
```

```
"compiler": "gc",
"platform": "linux/amd64"
}
```

- b. Verify that the cluster machine configuration API is accessible to the Machine config server resource, by running the following command and observing the output:

```
$ curl -v https://<loadbalancer_ip_address>:22623/healthz --insecure
```

If the configuration is correct, the output from the command shows the following response:

```
HTTP/1.1 200 OK
Content-Length: 0
```

- c. Verify that the controller is accessible to the Ingress Controller resource on port 80, by running the following command and observing the output:

```
$ curl -I -L -H "Host: console-openshift-console.apps.<cluster_name>.<base_domain>"
http://<load_balancer_front_end_IP_address>
```

If the configuration is correct, the output from the command shows the following response:

```
HTTP/1.1 302 Found
content-length: 0
location: https://console-openshift-console.apps.ocp4.private.opequon.net/
cache-control: no-cache
```

- d. Verify that the controller is accessible to the Ingress Controller resource on port 443, by running the following command and observing the output:

```
$ curl -I -L --insecure --resolve console-openshift-console.apps.<cluster_name>.<base_domain>:443:<Load Balancer Front End IP Address> https://console-openshift-console.apps.<cluster_name>.<base_domain>
```

If the configuration is correct, the output from the command shows the following response:

```
HTTP/1.1 200 OK
referrer-policy: strict-origin-when-cross-origin
set-cookie: csrf-token=UIYWOyQ62LWjw2h003xtYSKlh1a0Py2hhctw0WmV2YEdhJfYqWwCGBsja261dGLgaYO0nxzVERhiXt6QepA7g==; Path=/; Secure; SameSite=Lax
x-content-type-options: nosniff
x-dns-prefetch-control: off
x-frame-options: DENY
x-xss-protection: 1; mode=block
date: Wed, 04 Oct 2023 16:29:38 GMT
content-type: text/html; charset=utf-8
set-cookie: 1e2670d92730b515ce3a1bb65da45062=1bf5e9573c9a2760c964ed1659cc1673; path=/; HttpOnly; Secure; SameSite=None
cache-control: private
```

- Configure the DNS records for your cluster to target the front-end IP addresses of the user-managed load balancer. You must update records to your DNS server for the cluster API and applications over the load balancer.

Examples of modified DNS records

```
<load_balancer_ip_address> A api.<cluster_name>.<base_domain>
A record pointing to Load Balancer Front End
```

```
<load_balancer_ip_address> A apps.<cluster_name>.<base_domain>
A record pointing to Load Balancer Front End
```



IMPORTANT

DNS propagation might take some time for each DNS record to become available. Ensure that each DNS record propagates before validating each record.

- For your OpenShift Container Platform cluster to use the user-managed load balancer, you must specify the following configuration in your cluster's **install-config.yaml** file:

```
# ...
platform:
  vsphere:
    loadBalancer:
      type: UserManaged 1
      apiVIPs:
        - <api_ip> 2
      ingressVIPs:
        - <ingress_ip> 3
# ...
```

- Set **UserManaged** for the **type** parameter to specify a user-managed load balancer for your cluster. The parameter defaults to **OpenShiftManagedDefault**, which denotes the default internal load balancer. For services defined in an **openshift-kni-infra** namespace, a user-managed load balancer can deploy the **coredns** service to pods in your cluster but ignores **keepalived** and **haproxy** services.
- Required parameter when you specify a user-managed load balancer. Specify the user-managed load balancer's public IP address, so that the Kubernetes API can communicate with the user-managed load balancer.
- Required parameter when you specify a user-managed load balancer. Specify the user-managed load balancer's public IP address, so that the user-managed load balancer can manage ingress traffic for your cluster.

Verification

- Use the **curl** CLI command to verify that the user-managed load balancer and DNS record configuration are operational:
 - Verify that you can access the cluster API, by running the following command and observing the output:


```
$ curl https://api.<cluster_name>.<base_domain>:6443/version --insecure
```

If the configuration is correct, you receive a JSON object in response:

```
{
  "major": "1",
  "minor": "11+",
  "gitVersion": "v1.11.0+ad103ed",
  "gitCommit": "ad103ed",
  "gitTreeState": "clean",
  "buildDate": "2019-01-09T06:44:10Z",
  "goVersion": "go1.10.3",
  "compiler": "gc",
  "platform": "linux/amd64"
}
```

- b. Verify that you can access the cluster machine configuration, by running the following command and observing the output:

```
$ curl -v https://api.<cluster_name>.<base_domain>:22623/healthz --insecure
```

If the configuration is correct, the output from the command shows the following response:

```
HTTP/1.1 200 OK
Content-Length: 0
```

- c. Verify that you can access each cluster application on port, by running the following command and observing the output:

```
$ curl http://console-openshift-console.apps.<cluster_name>.<base_domain> -I -L --insecure
```

If the configuration is correct, the output from the command shows the following response:

```
HTTP/1.1 302 Found
content-length: 0
location: https://console-openshift-console.apps.<cluster-name>.<base domain>/
cache-control: no-cacheHTTP/1.1 200 OK
referrer-policy: strict-origin-when-cross-origin
set-cookie: csrf-
token=39HoZgztDnzjJkq/JuLJMeoKNXIfiVv2YgZc09c3TBOBU4NI6kDXaJH1LdicNhN1UsQ
Wzon4Dor9GWGfopaTEQ==; Path=/; Secure
x-content-type-options: nosniff
x-dns-prefetch-control: off
x-frame-options: DENY
x-xss-protection: 1; mode=block
date: Tue, 17 Nov 2020 08:42:10 GMT
content-type: text/html; charset=utf-8
set-cookie:
1e2670d92730b515ce3a1bb65da45062=9b714eb87e93cf34853e87a92d6894be; path=/;
HttpOnly; Secure; SameSite=None
cache-control: private
```

- d. Verify that you can access each cluster application on port 443, by running the following command and observing the output:

```
$ curl https://console-openshift-console.apps.<cluster_name>.<base_domain> -I -L --insecure
```

If the configuration is correct, the output from the command shows the following response:

```
HTTP/1.1 200 OK
referrer-policy: strict-origin-when-cross-origin
set-cookie: csrf-
token=UIYW0yQ62LWjw2h003xtYSKlh1a0Py2hhctw0WmV2YEdhJjFyQwWcGBsja261dG
LgaYO0nxzVERhiXt6QepA7g==; Path=/; Secure; SameSite=Lax
x-content-type-options: nosniff
x-dns-prefetch-control: off
x-frame-options: DENY
x-xss-protection: 1; mode=block
date: Wed, 04 Oct 2023 16:29:38 GMT
content-type: text/html; charset=utf-8
set-cookie:
1e2670d92730b515ce3a1bb65da45062=1bf5e9573c9a2760c964ed1659cc1673; path=/;
HttpOnly; Secure; SameSite=None
cache-control: private
```

23.2.4.6. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.




IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.
- Optional: Before you create the cluster, configure an external load balancer in place of the default load balancer.



IMPORTANT

You do not need to specify API and Ingress static addresses for your installation program. If you choose this configuration, you must take additional actions to define network targets that accept an IP address from each referenced vSphere subnet. See the section "Configuring a user-managed load balancer".

Procedure

- Change to the directory that contains the installation program and initialize the cluster deployment:


```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

- 1 For **<installation_directory>**, specify the location of your customized **./install-config.yaml** file.
- 2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to **<installation_directory>/openshift_install.log**.




IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

23.2.4.7. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

23.2.4.8. Creating registry storage

After you install the cluster, you must create storage for the registry Operator.

23.2.4.8.1. Image registry removed during installation

On platforms that do not provide shareable object storage, the OpenShift Image Registry Operator bootstraps itself as **Removed**. This allows **openshift-installer** to complete installations on these platform types.

After installation, you must edit the Image Registry Operator configuration to switch the **managementState** from **Removed** to **Managed**. When this has completed, you must configure storage.

23.2.4.8.2. Image registry storage configuration

The Image Registry Operator is not initially available for platforms that do not provide default storage. After installation, you must configure your registry to use storage so that the Registry Operator is made available.

Instructions are shown for configuring a persistent volume, which is required for production clusters. Where applicable, instructions are shown for configuring an empty directory as the storage location, which is available for only non-production clusters.

Additional instructions are provided for allowing the image registry to use block storage types by using the **Recreate** rollout strategy during upgrades.

23.2.4.8.2.1. Configuring registry storage for VMware vSphere

As a cluster administrator, following installation you must configure your registry to use storage.

Prerequisites

- Cluster administrator permissions.
- A cluster on VMware vSphere.
- Persistent storage provisioned for your cluster, such as Red Hat OpenShift Data Foundation.



IMPORTANT

OpenShift Container Platform supports **ReadWriteOnce** access for image registry storage when you have only one replica. **ReadWriteOnce** access also requires that the registry uses the **Recreate** rollout strategy. To deploy an image registry that supports high availability with two or more replicas, **ReadWriteMany** access is required.

- Must have "100Gi" capacity.



IMPORTANT

Testing shows issues with using the NFS server on RHEL as storage backend for core services. This includes the OpenShift Container Registry and Quay, Prometheus for monitoring storage, and Elasticsearch for logging storage. Therefore, using RHEL NFS to back PVs used by core services is not recommended.

Other NFS implementations on the marketplace might not have these issues. Contact the individual NFS implementation vendor for more information on any testing that was possibly completed against these OpenShift Container Platform core components.

Procedure

1. To configure your registry to use storage, change the **spec.storage.pvc** in the **configs.imageregistry/cluster** resource.



NOTE


When you use shared storage, review your security settings to prevent outside access.

2. Verify that you do not have a registry pod:

```
$ oc get pod -n openshift-image-registry -l docker-registry=default
```

Example output

```
No resources found in openshift-image-registry namespace
```

**NOTE**

If you do have a registry pod in your output, you do not need to continue with this procedure.

3. Check the registry configuration:

```
$ oc edit configs.imageregistry.operator.openshift.io
```

Example output

```
storage:
  pvc:
    claim: 1
```

- 1** Leave the **claim** field blank to allow the automatic creation of an **image-registry-storage** persistent volume claim (PVC). The PVC is generated based on the default storage class. However, be aware that the default storage class might provide ReadWriteOnce (RWO) volumes, such as a RADOS Block Device (RBD), which can cause issues when you replicate to more than one replica.

4. Check the **clusteroperator** status:


```
$ oc get clusteroperator image-registry
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED
image-registry	4.7	True	False	False

23.2.4.8.2.2. Configuring block registry storage for VMware vSphere

To allow the image registry to use block storage types such as vSphere Virtual Machine Disk (VMDK) during upgrades as a cluster administrator, you can use the **Recreate** rollout strategy.

**IMPORTANT**

Block storage volumes are supported but not recommended for use with image registry on production clusters. An installation where the registry is configured on block storage is not highly available because the registry cannot have more than one replica.

Procedure

1. Enter the following command to set the image registry storage as a block storage type, patch the registry so that it uses the **Recreate** rollout strategy, and runs with only **1** replica:

```
$ oc patch config.imageregistry.operator.openshift.io/cluster --type=merge -p '{"spec": {"rolloutStrategy": "Recreate", "replicas": 1}}'
```

2. Provision the PV for the block storage device, and create a PVC for that volume. The requested block volume uses the ReadWriteOnce (RWO) access mode.
 - a. Create a **pvc.yaml** file with the following contents to define a VMware vSphere **PersistentVolumeClaim** object:

```
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: image-registry-storage ❶
  namespace: openshift-image-registry ❷
spec:
  accessModes:
    - ReadWriteOnce ❸
  resources:
    requests:
      storage: 100Gi ❹
```

- ❶ A unique name that represents the **PersistentVolumeClaim** object.
- ❷ The namespace for the **PersistentVolumeClaim** object, which is **openshift-image-registry**.
- ❸ The access mode of the persistent volume claim. With **ReadWriteOnce**, the volume can be mounted with read and write permissions by a single node.
- ❹ The size of the persistent volume claim.

- b. Enter the following command to create the **PersistentVolumeClaim** object from the file:

```
$ oc create -f pvc.yaml -n openshift-image-registry
```

3. Enter the following command to edit the registry configuration so that it references the correct PVC:

```
$ oc edit config.imageregistry.operator.openshift.io -o yaml
```

Example output

```
storage:
  pvc:
    claim: ❶
```

- ❶ By creating a custom PVC, you can leave the **claim** field blank for the default automatic creation of an **image-registry-storage** PVC.

For instructions about configuring registry storage so that it references the correct PVC, see [Configuring the registry for vSphere](#).

23.2.4.9. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service


23.2.4.10. Next steps

- [Customize your cluster](#).
- If necessary, you can [opt out of remote health reporting](#) .
- [Set up your registry and configure registry storage](#) .
- Optional: [View the events from the vSphere Problem Detector Operator](#) to determine if the cluster has permission or storage configuration issues.

23.2.5. Installing a cluster on vSphere with network customizations

In OpenShift Container Platform version 4.16, you can install a cluster on your VMware vSphere instance by using installer-provisioned infrastructure with customized network configuration options. By customizing your network configuration, your cluster can coexist with existing IP address allocations in your environment and integrate with existing MTU and VXLAN configurations. To customize the installation, you modify parameters in the **install-config.yaml** file before you install the cluster.

You must set most of the network configuration parameters during installation, and you can modify only **kubeProxy** configuration parameters in a running cluster.



NOTE

OpenShift Container Platform supports deploying a cluster to a single VMware vCenter only. Deploying a cluster with machines/machine sets on multiple vCenters is not supported.

23.2.5.1. Prerequisites

- You have completed the tasks in [Preparing to install a cluster using installer-provisioned infrastructure](#).
- You reviewed your VMware platform licenses. Red Hat does not place any restrictions on your VMware licenses, but some VMware infrastructure components require licensing.
- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).

- You provisioned [persistent storage](#) for your cluster. To deploy a private image registry, your storage must provide **ReadWriteMany** access modes.
- The OpenShift Container Platform installer requires access to port 443 on the vCenter and ESXi hosts. You verified that port 443 is accessible.
- If you use a firewall, confirm with the administrator that port 443 is accessible. Control plane nodes must be able to reach vCenter and ESXi hosts on port 443 for the installation to succeed.
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.



NOTE


Be sure to also review this site list if you are configuring a proxy.

23.2.5.2. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.




IMPORTANT

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

23.2.5.3. VMware vSphere region and zone enablement

You can deploy an OpenShift Container Platform cluster to multiple vSphere datacenters that run in a single VMware vCenter. Each datacenter can run multiple clusters. This configuration reduces the risk of a hardware failure or network outage that can cause your cluster to fail. To enable regions and zones, you must define multiple failure domains for your OpenShift Container Platform cluster.



IMPORTANT

The VMware vSphere region and zone enablement feature requires the vSphere Container Storage Interface (CSI) driver as the default storage driver in the cluster. As a result, the feature is only available on a newly installed cluster.

For a cluster that was upgraded from a previous release, you must enable CSI automatic migration for the cluster. You can then configure multiple regions and zones for the upgraded cluster.

The default installation configuration deploys a cluster to a single vSphere datacenter. If you want to deploy a cluster to multiple vSphere datacenters, you must create an installation configuration file that enables the region and zone feature.

The default **install-config.yaml** file includes **vcenters** and **failureDomains** fields, where you can specify multiple vSphere datacenters and clusters for your OpenShift Container Platform cluster. You can leave these fields blank if you want to install an OpenShift Container Platform cluster in a vSphere environment that consists of single datacenter.

The following list describes terms associated with defining zones and regions for your cluster:

- Failure domain: Establishes the relationships between a region and zone. You define a failure domain by using vCenter objects, such as a **datastore** object. A failure domain defines the vCenter location for OpenShift Container Platform cluster nodes.
- Region: Specifies a vCenter datacenter. You define a region by using a tag from the **openshift-region** tag category.
- Zone: Specifies a vCenter cluster. You define a zone by using a tag from the **openshift-zone** tag category.



NOTE

If you plan on specifying more than one failure domain in your **install-config.yaml** file, you must create tag categories, zone tags, and region tags in advance of creating the configuration file.

You must create a vCenter tag for each vCenter datacenter, which represents a region. Additionally, you must create a vCenter tag for each cluster that runs in a datacenter, which represents a zone. After you create the tags, you must attach each tag to their respective datacenters and clusters.

The following table outlines an example of the relationship among regions, zones, and tags for a configuration with multiple vSphere datacenters running in a single VMware vCenter.

Datacenter (region)	Cluster (zone)	Tags
us-east	us-east-1	us-east-1a
		us-east-1b
	us-east-2	us-east-2a
		us-east-2b
us-west	us-west-1	us-west-1a
		us-west-1b
	us-west-2	us-west-2a
		us-west-2b

Additional resources

- [Additional VMware vSphere configuration parameters](#)
- [Deprecated VMware vSphere configuration parameters](#)
- [vSphere automatic migration](#)
- [VMware vSphere CSI Driver Operator](#)

23.2.5.4. Creating the installation configuration file

You can customize the OpenShift Container Platform cluster you install on VMware vSphere.

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create the **install-config.yaml** file.
 - a. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install create install-config --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the directory name to store the files that the installation program creates.

When specifying the directory:

- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
 - Use an empty directory. Some installation assets, such as bootstrap X.509 certificates, have short expiration intervals, therefore you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.
- b. At the prompts, provide the configuration details for your cloud:
 - i. Optional: Select an SSH key to use to access your cluster machines.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- ii. Select **vsphere** as the platform to target.

- iii. Specify the name of your vCenter instance.
- iv. Specify the user name and password for the vCenter account that has the required permissions to create the cluster.
The installation program connects to your vCenter instance.
- v. Select the data center in your vCenter instance to connect to.



NOTE

After you create the installation configuration file, you can modify the file to create a multiple vSphere datacenters environment. This means that you can deploy an OpenShift Container Platform cluster to multiple vSphere datacenters that run in a single VMware vCenter. For more information about creating this environment, see the section named *VMware vSphere region and zone enablement*.

- vi. Select the default vCenter datastore to use.



WARNING


You can specify the path of any datastore that exists in a datastore cluster. By default, Storage Distributed Resource Scheduler (SDRS), which uses Storage vMotion, is automatically enabled for a datastore cluster. Red Hat does not support Storage vMotion, so you must disable Storage DRS to avoid data loss issues for your OpenShift Container Platform cluster.

You cannot specify more than one datastore path. If you must specify VMs across multiple datastores, use a **datastore** object to specify a failure domain in your cluster's **install-config.yaml** configuration file. For more information, see "VMware vSphere region and zone enablement".

- vii. Select the vCenter cluster to install the OpenShift Container Platform cluster in. The installation program uses the root resource pool of the vSphere cluster as the default resource pool.
- viii. Select the network in the vCenter instance that contains the virtual IP addresses and DNS records that you configured.
- ix. Enter the virtual IP address that you configured for control plane API access.
- x. Enter the virtual IP address that you configured for cluster ingress.
- xi. Enter the base domain. This base domain must be the same one that you used in the DNS records that you configured.
- xii. Enter a descriptive name for your cluster.

The cluster name you enter must match the cluster name you specified when configuring the DNS records.

2. Modify the **install-config.yaml** file. You can find more information about the available parameters in the "Installation configuration parameters" section.
3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

Additional resources

- [Installation configuration parameters](#)

23.2.5.4.1. Sample install-config.yaml file for an installer-provisioned VMware vSphere cluster

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.

```

apiVersion: v1
baseDomain: example.com 1
compute: 2
- architecture: amd64
  name: <worker_node>
  platform: {}
  replicas: 3
controlPlane: 3
  architecture: amd64
  name: <parent_node>
  platform: {}
  replicas: 3
metadata:
  creationTimestamp: null
  name: test 4
networking:
  clusterNetwork:
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  machineNetwork:
  - cidr: 10.0.0.0/16
  networkType: OVNKubernetes 5
  serviceNetwork:
  - 172.30.0.0/16
platform:
  vsphere: 6
  apiVIPs:
  - 10.0.0.1
  failureDomains: 7
  - name: <failure_domain_name>
    region: <default_region_name>
    server: <fully_qualified_domain_name>

```

```

topology:
  computeCluster: "/<datacenter>/host/<cluster>"
  datacenter: <datacenter>
  datastore: "/<datacenter>/datastore/<datastore>" 8
  networks:
  - <VM_Network_name>
  resourcePool: "/<datacenter>/host/<cluster>/Resources/<resourcePool>" 9
  folder: "/<datacenter_name>/vm/<folder_name>/<subfolder_name>"
  tagIDs: 10
  - <tag_id> 11
  zone: <default_zone_name>
ingressVIPs:
- 10.0.0.2
vcenters:
- datacenters:
- <datacenter>
password: <password>
port: 443
server: <fully_qualified_domain_name>
user: administrator@vsphere.local
diskType: thin 12
fips: false
pullSecret: '{"auths": ...}'
sshKey: 'ssh-ed25519 AAAA...'

```

- 1 The base domain of the cluster. All DNS records must be sub-domains of this base and include the cluster name.
- 2 3 The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.
- 4 The cluster name that you specified in your DNS records.
- 6 Optional: Provides additional configuration for the machine pool parameters for the compute and control plane machines.
- 7 Establishes the relationships between a region and zone. You define a failure domain by using vCenter objects, such as a **datastore** object. A failure domain defines the vCenter location for OpenShift Container Platform cluster nodes.
- 8 The path to the vSphere datastore that holds virtual machine files, templates, and ISO images.




IMPORTANT

You can specify the path of any datastore that exists in a datastore cluster. By default, Storage vMotion is automatically enabled for a datastore cluster. Red Hat does not support Storage vMotion, so you must disable Storage vMotion to avoid data loss issues for your OpenShift Container Platform cluster.

If you must specify VMs across multiple datastores, use a **datastore** object to specify a failure domain in your cluster's **install-config.yaml** configuration file. For more information, see "VMware vSphere region and zone enablement".

- 9 Optional: Provides an existing resource pool for machine creation. If you do not specify a value, the installation program uses the root resource pool of the vSphere cluster.
- 10 Optional: Each VM created by OpenShift Container Platform is assigned a unique tag that is specific to the cluster. The assigned tag enables the installation program to identify and remove the associated VMs when a cluster is decommissioned. You can list up to ten additional tag IDs to be attached to the VMs provisioned by the installation program.
- 11 The ID of the tag to be associated by the installation program. For example, **urn:vmomi:InventoryServiceTag:208e713c-cae3-4b7f-918e-4051ca7d1f97:GLOBAL**. For more information about determining the tag ID, see the [vSphere Tags and Attributes documentation](#).
- 12 The vSphere disk provisioning method.
- 5 The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.



NOTE


In OpenShift Container Platform 4.12 and later, the **apiVIP** and **ingressVIP** configuration settings are deprecated. Instead, use a list format to enter values in the **apiVIPs** and **ingressVIPs** configuration settings.

23.2.5.4.2. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
```

```

baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5

```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations. You must include vCenter's IP address and the IP range that you use for its machines.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.



NOTE

The installation program does not support the proxy **readinessEndpoints** field.




NOTE

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.

**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.


23.2.5.4.3. Optional: Deploying with dual-stack networking

For dual-stack networking in OpenShift Container Platform clusters, you can configure IPv4 and IPv6 address endpoints for cluster nodes. To configure IPv4 and IPv6 address endpoints for cluster nodes, edit the **machineNetwork**, **clusterNetwork**, and **serviceNetwork** configuration settings in the **install-config.yaml** file. Each setting must have two CIDR entries each. For a cluster with the IPv4 family as the primary address family, specify the IPv4 setting first. For a cluster with the IPv6 family as the primary address family, specify the IPv6 setting first.

```
machineNetwork:
- cidr: {{ extcidrnet }}
- cidr: {{ extcidrnet6 }}
clusterNetwork:
- cidr: 10.128.0.0/14
  hostPrefix: 23
- cidr: fd02::/48
  hostPrefix: 64
serviceNetwork:
- 172.30.0.0/16
- fd03::/112
```

To provide an interface to the cluster for applications that use IPv4 and IPv6 addresses, configure IPv4 and IPv6 virtual IP (VIP) address endpoints for the Ingress VIP and API VIP services. To configure IPv4 and IPv6 address endpoints, edit the **apiVIPs** and **ingressVIPs** configuration settings in the **install-config.yaml** file. The **apiVIPs** and **ingressVIPs** configuration settings use a list format. The order of the list indicates the primary and secondary VIP address for each service.

```
platform:
vsphere:
  apiVIPs:
- <api_ipv4>
- <api_ipv6>
  ingressVIPs:
- <wildcard_ipv4>
- <wildcard_ipv6>
```

**NOTE**

For a cluster with dual-stack networking configuration, you must assign both IPv4 and IPv6 addresses to the same interface.

23.2.5.4.4. Configuring regions and zones for a VMware vCenter

You can modify the default installation configuration file, so that you can deploy an OpenShift Container Platform cluster to multiple vSphere datacenters that run in a single VMware vCenter.

The default **install-config.yaml** file configuration from the previous release of OpenShift Container Platform is deprecated. You can continue to use the deprecated default configuration, but the **openshift-installer** will prompt you with a warning message that indicates the use of deprecated fields

in the configuration file.



IMPORTANT

The example uses the **govc** command. The **govc** command is an open source command available from VMware; it is not available from Red Hat. The Red Hat support team does not maintain the **govc** command. Instructions for downloading and installing **govc** are found on the VMware documentation website

Prerequisites

- You have an existing **install-config.yaml** installation configuration file.



IMPORTANT

You must specify at least one failure domain for your OpenShift Container Platform cluster, so that you can provision datacenter objects for your VMware vCenter server. Consider specifying multiple failure domains if you need to provision virtual machine nodes in different datacenters, clusters, datastores, and other components. To enable regions and zones, you must define multiple failure domains for your OpenShift Container Platform cluster.

Procedure

1. Enter the following **govc** command-line tool commands to create the **openshift-region** and **openshift-zone** vCenter tag categories:



IMPORTANT

If you specify different names for the **openshift-region** and **openshift-zone** vCenter tag categories, the installation of the OpenShift Container Platform cluster fails.

```
$ govc tags.category.create -d "OpenShift region" openshift-region
```

```
$ govc tags.category.create -d "OpenShift zone" openshift-zone
```

2. To create a region tag for each region vSphere datacenter where you want to deploy your cluster, enter the following command in your terminal:

```
$ govc tags.create -c <region_tag_category> <region_tag>
```

3. To create a zone tag for each vSphere cluster where you want to deploy your cluster, enter the following command:

```
$ govc tags.create -c <zone_tag_category> <zone_tag>
```

4. Attach region tags to each vCenter datacenter object by entering the following command:

```
$ govc tags.attach -c <region_tag_category> <region_tag_1> /<datacenter_1>
```

5. Attach the zone tags to each vCenter datacenter object by entering the following command:

```
$ govc tags.attach -c <zone_tag_category> <zone_tag_1> /<datacenter_1>/host/vcs-mdcnc-
workload-1
```

6. Change to the directory that contains the installation program and initialize the cluster deployment according to your chosen installation requirements.

Sample install-config.yaml file with multiple datacenters defined in a vSphere center

```
---
compute:
---
vsphere:
  zones:
    - "<machine_pool_zone_1>"
    - "<machine_pool_zone_2>"
---
controlPlane:
---
vsphere:
  zones:
    - "<machine_pool_zone_1>"
    - "<machine_pool_zone_2>"
---
platform:
  vsphere:
    vcenters:
---
  datacenters:
    - <datacenter1_name>
    - <datacenter2_name>
  failureDomains:
    - name: <machine_pool_zone_1>
      region: <region_tag_1>
      zone: <zone_tag_1>
      server: <fully_qualified_domain_name>
      topology:
        datacenter: <datacenter1>
        computeCluster: "/<datacenter1>/host/<cluster1>"
        networks:
          - <VM_Network1_name>
        datastore: "/<datacenter1>/datastore/<datastore1>"
        resourcePool: "/<datacenter1>/host/<cluster1>/Resources/<resourcePool1>"
        folder: "/<datacenter1>/vm/<folder1>"
    - name: <machine_pool_zone_2>
      region: <region_tag_2>
      zone: <zone_tag_2>
      server: <fully_qualified_domain_name>
      topology:
        datacenter: <datacenter2>
        computeCluster: "/<datacenter2>/host/<cluster2>"
        networks:
          - <VM_Network2_name>
        datastore: "/<datacenter2>/datastore/<datastore2>"
```

```
resourcePool: "/<datacenter2>/host/<cluster2>/Resources/<resourcePool2>"
folder: "/<datacenter2>/vm/<folder2>"
```

23.2.5.5. Network configuration phases

There are two phases prior to OpenShift Container Platform installation where you can customize the network configuration.

Phase 1

You can customize the following network-related fields in the **install-config.yaml** file before you create the manifest files:

- **networking.networkType**
- **networking.clusterNetwork**
- **networking.serviceNetwork**
- **networking.machineNetwork**

For more information on these fields, refer to *Installation configuration parameters*.



NOTE

Set the **networking.machineNetwork** to match the CIDR that the preferred NIC resides in.



IMPORTANT

The CIDR range **172.17.0.0/16** is reserved by libVirt. You cannot use this range or any range that overlaps with this range for any networks in your cluster.

Phase 2

After creating the manifest files by running **openshift-install create manifests**, you can define a customized Cluster Network Operator manifest with only the fields you want to modify. You can use the manifest to specify advanced network configuration.

You cannot override the values specified in phase 1 in the **install-config.yaml** file during phase 2. However, you can further customize the network plugin during phase 2.

23.2.5.6. Specifying advanced network configuration

You can use advanced network configuration for your network plugin to integrate your cluster into your existing network environment. You can specify advanced network configuration only before you install the cluster.



IMPORTANT

Customizing your network configuration by modifying the OpenShift Container Platform manifest files created by the installation program is not supported. Applying a manifest file that you create, as in the following procedure, is supported.

Prerequisites

- You have created the **install-config.yaml** file and completed any modifications to it.

Procedure

- Change to the directory that contains the installation program and create the manifests:

```
$ ./openshift-install create manifests --dir <installation_directory> 1
```

- 1** **<installation_directory>** specifies the name of the directory that contains the **install-config.yaml** file for your cluster.

- Create a stub manifest file for the advanced network configuration that is named **cluster-network-03-config.yml** in the **<installation_directory>/manifests/** directory:

```
apiVersion: operator.openshift.io/v1
kind: Network
metadata:
  name: cluster
spec:
```

- Specify the advanced network configuration for your cluster in the **cluster-network-03-config.yml** file, such as in the following example:

Enable IPsec for the OVN-Kubernetes network provider

```
apiVersion: operator.openshift.io/v1
kind: Network
metadata:
  name: cluster
spec:
  defaultNetwork:
    ovnKubernetesConfig:
      ipsecConfig:
        mode: Full
```

- Optional: Back up the **manifests/cluster-network-03-config.yml** file. The installation program consumes the **manifests/** directory when you create the Ignition config files.
- Remove the Kubernetes manifest files that define the control plane machines and compute machineSets:

```
$ rm -f openshift/99_openshift-cluster-api_master-machines-*.yaml openshift/99_openshift-cluster-api_worker-machineset-*.yaml
```

Because you create and manage these resources yourself, you do not have to initialize them.

- You can preserve the MachineSet files to create compute machines by using the machine API, but you must update references to them to match your environment.

23.2.5.7. Cluster Network Operator configuration

The configuration for the cluster network is specified as part of the Cluster Network Operator (CNO) configuration and stored in a custom resource (CR) object that is named **cluster**. The CR specifies the fields for the **Network** API in the **operator.openshift.io** API group.

The CNO configuration inherits the following fields during cluster installation from the **Network** API in the **Network.config.openshift.io** API group:

clusterNetwork

IP address pools from which pod IP addresses are allocated.

serviceNetwork

IP address pool for services.

defaultNetwork.type

Cluster network plugin. **OVNKubernetes** is the only supported plugin during installation.

You can specify the cluster network plugin configuration for your cluster by setting the fields for the **defaultNetwork** object in the CNO object named **cluster**.

23.2.5.7.1. Cluster Network Operator configuration object

The fields for the Cluster Network Operator (CNO) are described in the following table:

Table 23.7. Cluster Network Operator configuration object


Field	Type	Description
metadata.name	string	The name of the CNO object. This name is always cluster .
spec.clusterNetwork	array	<p>A list specifying the blocks of IP addresses from which pod IP addresses are allocated and the subnet prefix length assigned to each individual node in the cluster. For example:</p> <pre>spec: clusterNetwork: - cidr: 10.128.0.0/19 hostPrefix: 23 - cidr: 10.128.32.0/19 hostPrefix: 23</pre>
spec.serviceNetwork	array	<p>A block of IP addresses for services. The OpenShift SDN and OVN-Kubernetes network plugins support only a single IP address block for the service network. For example:</p> <pre>spec: serviceNetwork: - 172.30.0.0/14</pre> <p>You can customize this field only in the install-config.yaml file before you create the manifests. The value is read-only in the manifest file.</p>

Field	Type	Description
spec.defaultNetwork	object	Configures the network plugin for the cluster network.
spec.kubeProxyConfig	object	The fields for this object specify the kube-proxy configuration. If you are using the OVN-Kubernetes cluster network plugin, the kube-proxy configuration has no effect.

defaultNetwork object configuration

The values for the **defaultNetwork** object are defined in the following table:

Table 23.8. **defaultNetwork** object

Field	Type	Description
type	string	<p>OVNKubernetes. The Red Hat OpenShift Networking network plugin is selected during installation. This value cannot be changed after cluster installation.</p> <div style="display: flex; align-items: flex-start;">  <div> <p>NOTE</p> <p>OpenShift Container Platform uses the OVN-Kubernetes network plugin by default. OpenShift SDN is no longer available as an installation choice for new clusters.</p> </div> </div>
ovnKubernetesConfig	object	This object is only valid for the OVN-Kubernetes network plugin.

Configuration for the OVN-Kubernetes network plugin

The following table describes the configuration fields for the OVN-Kubernetes network plugin:

Table 23.9. **ovnKubernetesConfig** object

Field	Type	Description
-------	------	-------------


Field	Type	Description
mtu	integer	<p>The maximum transmission unit (MTU) for the Geneve (Generic Network Virtualization Encapsulation) overlay network. This is detected automatically based on the MTU of the primary network interface. You do not normally need to override the detected MTU.</p> <p>If the auto-detected value is not what you expect it to be, confirm that the MTU on the primary network interface on your nodes is correct. You cannot use this option to change the MTU value of the primary network interface on the nodes.</p> <p>If your cluster requires different MTU values for different nodes, you must set this value to 100 less than the lowest MTU value in your cluster. For example, if some nodes in your cluster have an MTU of 9001, and some have an MTU of 1500, you must set this value to 1400.</p>
genevePort	integer	The port to use for all Geneve packets. The default value is 6081 . This value cannot be changed after cluster installation.
ipsecConfig	object	Specify a configuration object for customizing the IPsec configuration.
ipv4	object	Specifies a configuration object for IPv4 settings.
ipv6	object	Specifies a configuration object for IPv6 settings.
policyAuditConfig	object	Specify a configuration object for customizing network policy audit logging. If unset, the defaults audit log settings are used.
gatewayConfig	object	<p>Optional: Specify a configuration object for customizing how egress traffic is sent to the node gateway.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p>NOTE</p> <p>While migrating egress traffic, you can expect some disruption to workloads and service traffic until the Cluster Network Operator (CNO) successfully rolls out the changes.</p> </div> </div>

Table 23.10. `ovnKubernetesConfig.ipv4` object

Field	Type	Description
-------	------	-------------

Field	Type	Description
internalTransitSwitchSubnet	string	<p>If your existing network infrastructure overlaps with the 100.88.0.0/16 IPv4 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. The subnet for the distributed transit switch that enables east-west traffic. This subnet cannot overlap with any other subnets used by OVN-Kubernetes or on the host itself. It must be large enough to accommodate one IP address per node in your cluster.</p> <p>The default value is 100.88.0.0/16.</p>
internalJoinSubnet	string	<p>If your existing network infrastructure overlaps with the 100.64.0.0/16 IPv4 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. You must ensure that the IP address range does not overlap with any other subnet used by your OpenShift Container Platform installation. The IP address range must be larger than the maximum number of nodes that can be added to the cluster. For example, if the clusterNetwork.cidr value is 10.128.0.0/14 and the clusterNetwork.hostPrefix value is /23, then the maximum number of nodes is 2⁽²³⁻¹⁴⁾=512.</p> <p>The default value is 100.64.0.0/16.</p>

Table 23.11. `ovnKubernetesConfig.ipv6` object

Field	Type	Description
internalTransitSwitchSubnet	string	<p>If your existing network infrastructure overlaps with the fd97::<64 IPv6 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. The subnet for the distributed transit switch that enables east-west traffic. This subnet cannot overlap with any other subnets used by OVN-Kubernetes or on the host itself. It must be large enough to accommodate one IP address per node in your cluster.</p> <p>The default value is fd97::<64.</p>
internalJoinSubnet	string	<p>If your existing network infrastructure overlaps with the fd98::<64 IPv6 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. You must ensure that the IP address range does not overlap with any other subnet used by your OpenShift Container Platform installation. The IP address range must be larger than the maximum number of nodes that can be added to the cluster.</p> <p>The default value is fd98::<64.</p>

Table 23.12. `policyAuditConfig` object

Field	Type	Description
rateLimit	integer	The maximum number of messages to generate every second per node. The default value is 20 messages per second.
maxFileSize	integer	The maximum size for the audit log in bytes. The default value is 50000000 or 50 MB.
maxLogFiles	integer	The maximum number of log files that are retained.
destination	string	One of the following additional audit log targets: libc The libc syslog() function of the journald process on the host. udp:<host>:<port> A syslog server. Replace <host>:<port> with the host and port of the syslog server. unix:<file> A Unix Domain Socket file specified by <file> . null Do not send the audit logs to any additional target.
syslogFacility	string	The syslog facility, such as kern , as defined by RFC5424. The default value is local0 .

Table 23.13. gatewayConfig object

Field	Type	Description
routingViaHost	boolean	Set this field to true to send egress traffic from pods to the host networking stack. For highly-specialized installations and applications that rely on manually configured routes in the kernel routing table, you might want to route egress traffic to the host networking stack. By default, egress traffic is processed in OVN to exit the cluster and is not affected by specialized routes in the kernel routing table. The default value is false . This field has an interaction with the Open vSwitch hardware offloading feature. If you set this field to true , you do not receive the performance benefits of the offloading because egress traffic is processed by the host networking stack.
ipForwarding	object	You can control IP forwarding for all traffic on OVN-Kubernetes managed interfaces by using the ipForwarding specification in the Network resource. Specify Restricted to only allow IP forwarding for Kubernetes related traffic. Specify Global to allow forwarding of all IP traffic. For new installations, the default is Restricted . For updates to OpenShift Container Platform 4.14 or later, the default is Global .

Field	Type	Description
ipv4	object	Optional: Specify an object to configure the internal OVN-Kubernetes masquerade address for host to service traffic for IPv4 addresses.
ipv6	object	Optional: Specify an object to configure the internal OVN-Kubernetes masquerade address for host to service traffic for IPv6 addresses.

Table 23.14. gatewayConfig.ipv4 object

Field	Type	Description
internalMasqueradeSubnet	string	The masquerade IPv4 addresses that are used internally to enable host to service traffic. The host is configured with these IP addresses as well as the shared gateway bridge interface. The default value is 169.254.169.0/29 .

Table 23.15. gatewayConfig.ipv6 object

Field	Type	Description
internalMasqueradeSubnet	string	The masquerade IPv6 addresses that are used internally to enable host to service traffic. The host is configured with these IP addresses as well as the shared gateway bridge interface. The default value is fd69::/125 .

Table 23.16. ipsecConfig object

Field	Type	Description
mode	string	Specifies the behavior of the IPsec implementation. Must be one of the following values: <ul style="list-style-type: none"> ● Disabled: IPsec is not enabled on cluster nodes. ● External: IPsec is enabled for network traffic with external hosts. ● Full: IPsec is enabled for pod traffic and network traffic with external hosts.

Example OVN-Kubernetes configuration with IPsec enabled

```
defaultNetwork:
  type: OVNKubernetes
  ovnKubernetesConfig:
    mtu: 1400
```

```
genevePort: 6081
ipsecConfig:
mode: Full
```


**IMPORTANT**

Using OVNKubernetes can lead to a stack exhaustion problem on IBM Power®.

kubeProxyConfig object configuration (OpenShiftSDN container network interface only)

The values for the **kubeProxyConfig** object are defined in the following table:

Table 23.17. **kubeProxyConfig** object

Field	Type	Description
iptablesSyncPeriod	string	<p>The refresh period for iptables rules. The default value is 30s. Valid suffixes include s, m, and h and are described in the Go time package documentation.</p> <div style="display: flex; align-items: flex-start;">  <div> <p>NOTE</p> <p>Because of performance improvements introduced in OpenShift Container Platform 4.3 and greater, adjusting the iptablesSyncPeriod parameter is no longer necessary.</p> </div> </div>
proxyArguments.iptables-min-sync-period	array	<p>The minimum duration before refreshing iptables rules. This field ensures that the refresh does not happen too frequently. Valid suffixes include s, m, and h and are described in the Go time package. The default value is:</p> <pre>kubeProxyConfig: proxyArguments: iptables-min-sync-period: - 0s</pre>

23.2.5.8. Services for a user-managed load balancer

You can configure an OpenShift Container Platform cluster to use a user-managed load balancer in place of the default load balancer.

**IMPORTANT**

Configuring a user-managed load balancer depends on your vendor's load balancer.


The information and examples in this section are for guideline purposes only. Consult the vendor documentation for more specific information about the vendor's load balancer.

Red Hat supports the following services for a user-managed load balancer:

- Ingress Controller
- OpenShift API
- OpenShift MachineConfig API

You can choose whether you want to configure one or all of these services for a user-managed load balancer. Configuring only the Ingress Controller service is a common configuration option. To better understand each service, view the following diagrams:

Figure 23.4. Example network workflow that shows an Ingress Controller operating in an OpenShift Container Platform environment



496_OpenShift_1223

Figure 23.5. Example network workflow that shows an OpenShift API operating in an OpenShift Container Platform environment





Figure 23.6. Example network workflow that shows an OpenShift MachineConfig API operating in an OpenShift Container Platform environment



The following configuration options are supported for user-managed load balancers:

- Use a node selector to map the Ingress Controller to a specific set of nodes. You must assign a static IP address to each node in this set, or configure each node to receive the same IP address from the Dynamic Host Configuration Protocol (DHCP). Infrastructure nodes commonly receive this type of configuration.

- Target all IP addresses on a subnet. This configuration can reduce maintenance overhead, because you can create and destroy nodes within those networks without reconfiguring the load balancer targets. If you deploy your ingress pods by using a machine set on a smaller network, such as a `/27` or `/28`, you can simplify your load balancer targets.

TIP


You can list all IP addresses that exist in a network by checking the machine config pool's resources.

Before you configure a user-managed load balancer for your OpenShift Container Platform cluster, consider the following information:

- For a front-end IP address, you can use the same IP address for the front-end IP address, the Ingress Controller's load balancer, and API load balancer. Check the vendor's documentation for this capability.
- For a back-end IP address, ensure that an IP address for an OpenShift Container Platform control plane node does not change during the lifetime of the user-managed load balancer. You can achieve this by completing one of the following actions:
 - Assign a static IP address to each control plane node.
 - Configure each node to receive the same IP address from the DHCP every time the node requests a DHCP lease. Depending on the vendor, the DHCP lease might be in the form of an IP reservation or a static DHCP assignment.
- Manually define each node that runs the Ingress Controller in the user-managed load balancer for the Ingress Controller back-end service. For example, if the Ingress Controller moves to an undefined node, a connection outage can occur.

23.2.5.8.1. Configuring a user-managed load balancer

You can configure an OpenShift Container Platform cluster to use a user-managed load balancer in place of the default load balancer.



IMPORTANT

Before you configure a user-managed load balancer, ensure that you read the "Services for a user-managed load balancer" section.

Read the following prerequisites that apply to the service that you want to configure for your user-managed load balancer.



NOTE

MetalLB, which runs on a cluster, functions as a user-managed load balancer.

OpenShift API prerequisites

- You defined a front-end IP address.
- TCP ports 6443 and 22623 are exposed on the front-end IP address of your load balancer. Check the following items:

- Port 6443 provides access to the OpenShift API service.
- Port 22623 can provide ignition startup configurations to nodes.
- The front-end IP address and port 6443 are reachable by all users of your system with a location external to your OpenShift Container Platform cluster.
- The front-end IP address and port 22623 are reachable only by OpenShift Container Platform nodes.
- The load balancer backend can communicate with OpenShift Container Platform control plane nodes on port 6443 and 22623.

Ingress Controller prerequisites

- You defined a front-end IP address.
- TCP ports 443 and 80 are exposed on the front-end IP address of your load balancer.
- The front-end IP address, port 80 and port 443 are be reachable by all users of your system with a location external to your OpenShift Container Platform cluster.
- The front-end IP address, port 80 and port 443 are reachable to all nodes that operate in your OpenShift Container Platform cluster.
- The load balancer backend can communicate with OpenShift Container Platform nodes that run the Ingress Controller on ports 80, 443, and 1936.

Prerequisite for health check URL specifications

You can configure most load balancers by setting health check URLs that determine if a service is available or unavailable. OpenShift Container Platform provides these health checks for the OpenShift API, Machine Configuration API, and Ingress Controller backend services.

The following examples show health check specifications for the previously listed backend services:

Example of a Kubernetes API health check specification

```
Path: HTTPS:6443/readyz
Healthy threshold: 2
Unhealthy threshold: 2
Timeout: 10
Interval: 10
```

Example of a Machine Config API health check specification

```
Path: HTTPS:22623/healthz
Healthy threshold: 2
Unhealthy threshold: 2
Timeout: 10
Interval: 10
```

Example of an Ingress Controller health check specification

```
Path: HTTP:1936/healthz/ready
Healthy threshold: 2
```


Unhealthy threshold: 2
 Timeout: 5
 Interval: 10

Procedure

1. Configure the HAProxy Ingress Controller, so that you can enable access to the cluster from your load balancer on ports 6443, 22623, 443, and 80. Depending on your needs, you can specify the IP address of a single subnet or IP addresses from multiple subnets in your HAProxy configuration.

Example HAProxy configuration with one listed subnet

```
# ...
listen my-cluster-api-6443
  bind 192.168.1.100:6443
  mode tcp
  balance roundrobin
  option httpchk
  http-check connect
  http-check send meth GET uri /readyz
  http-check expect status 200
  server my-cluster-master-2 192.168.1.101:6443 check inter 10s rise 2 fall 2
  server my-cluster-master-0 192.168.1.102:6443 check inter 10s rise 2 fall 2
  server my-cluster-master-1 192.168.1.103:6443 check inter 10s rise 2 fall 2

listen my-cluster-machine-config-api-22623
  bind 192.168.1.100:22623
  mode tcp
  balance roundrobin
  option httpchk
  http-check connect
  http-check send meth GET uri /healthz
  http-check expect status 200
  server my-cluster-master-2 192.168.1.101:22623 check inter 10s rise 2 fall 2
  server my-cluster-master-0 192.168.1.102:22623 check inter 10s rise 2 fall 2
  server my-cluster-master-1 192.168.1.103:22623 check inter 10s rise 2 fall 2

listen my-cluster-apps-443
  bind 192.168.1.100:443
  mode tcp
  balance roundrobin
  option httpchk
  http-check connect
  http-check send meth GET uri /healthz/ready
  http-check expect status 200
  server my-cluster-worker-0 192.168.1.111:443 check port 1936 inter 10s rise 2 fall 2
  server my-cluster-worker-1 192.168.1.112:443 check port 1936 inter 10s rise 2 fall 2
  server my-cluster-worker-2 192.168.1.113:443 check port 1936 inter 10s rise 2 fall 2

listen my-cluster-apps-80
  bind 192.168.1.100:80
  mode tcp
  balance roundrobin
  option httpchk
```

```
http-check connect
http-check send meth GET uri /healthz/ready
http-check expect status 200
server my-cluster-worker-0 192.168.1.111:80 check port 1936 inter 10s rise 2 fall 2
server my-cluster-worker-1 192.168.1.112:80 check port 1936 inter 10s rise 2 fall 2
server my-cluster-worker-2 192.168.1.113:80 check port 1936 inter 10s rise 2 fall 2
# ...
```

Example HAProxy configuration with multiple listed subnets

```
# ...
listen api-server-6443
bind *:6443
mode tcp
server master-00 192.168.83.89:6443 check inter 1s
server master-01 192.168.84.90:6443 check inter 1s
server master-02 192.168.85.99:6443 check inter 1s
server bootstrap 192.168.80.89:6443 check inter 1s

listen machine-config-server-22623
bind *:22623
mode tcp
server master-00 192.168.83.89:22623 check inter 1s
server master-01 192.168.84.90:22623 check inter 1s
server master-02 192.168.85.99:22623 check inter 1s
server bootstrap 192.168.80.89:22623 check inter 1s

listen ingress-router-80
bind *:80
mode tcp
balance source
server worker-00 192.168.83.100:80 check inter 1s
server worker-01 192.168.83.101:80 check inter 1s

listen ingress-router-443
bind *:443
mode tcp
balance source
server worker-00 192.168.83.100:443 check inter 1s
server worker-01 192.168.83.101:443 check inter 1s

listen ironic-api-6385
bind *:6385
mode tcp
balance source
server master-00 192.168.83.89:6385 check inter 1s
server master-01 192.168.84.90:6385 check inter 1s
server master-02 192.168.85.99:6385 check inter 1s
server bootstrap 192.168.80.89:6385 check inter 1s

listen inspector-api-5050
bind *:5050
mode tcp
balance source
server master-00 192.168.83.89:5050 check inter 1s
server master-01 192.168.84.90:5050 check inter 1s
```

```
server master-02 192.168.85.99:5050 check inter 1s
server bootstrap 192.168.80.89:5050 check inter 1s
# ...
```

2. Use the **curl** CLI command to verify that the user-managed load balancer and its resources are operational:
 - a. Verify that the cluster machine configuration API is accessible to the Kubernetes API server resource, by running the following command and observing the response:

```
$ curl https://<loadbalancer_ip_address>:6443/version --insecure
```

If the configuration is correct, you receive a JSON object in response:

```
{
  "major": "1",
  "minor": "11+",
  "gitVersion": "v1.11.0+ad103ed",
  "gitCommit": "ad103ed",
  "gitTreeState": "clean",
  "buildDate": "2019-01-09T06:44:10Z",
  "goVersion": "go1.10.3",
  "compiler": "gc",
  "platform": "linux/amd64"
}
```

- b. Verify that the cluster machine configuration API is accessible to the Machine config server resource, by running the following command and observing the output:

```
$ curl -v https://<loadbalancer_ip_address>:22623/healthz --insecure
```

If the configuration is correct, the output from the command shows the following response:

```
HTTP/1.1 200 OK
Content-Length: 0
```

- c. Verify that the controller is accessible to the Ingress Controller resource on port 80, by running the following command and observing the output:

```
$ curl -I -L -H "Host: console-openshift-console.apps.<cluster_name>.<base_domain>"
http://<load_balancer_front_end_IP_address>
```

If the configuration is correct, the output from the command shows the following response:

```
HTTP/1.1 302 Found
content-length: 0
location: https://console-openshift-console.apps.ocp4.private.opequon.net/
cache-control: no-cache
```

- d. Verify that the controller is accessible to the Ingress Controller resource on port 443, by running the following command and observing the output:

```
$ curl -I -L --insecure --resolve console-openshift-console.apps.<cluster_name>.<base_domain>:443:<Load Balancer Front End IP Address> https://console-openshift-console.apps.<cluster_name>.<base_domain>
```

If the configuration is correct, the output from the command shows the following response:


```
HTTP/1.1 200 OK
referrer-policy: strict-origin-when-cross-origin
set-cookie: csrf-token=UIYW0yQ62LWjw2h003xtYSKlh1a0Py2hhctw0WmV2YEdhJfYqWwCGBsja261dGLgaYO0nxzVERhiXt6QepA7g==; Path=/; Secure; SameSite=Lax
x-content-type-options: nosniff
x-dns-prefetch-control: off
x-frame-options: DENY
x-xss-protection: 1; mode=block
date: Wed, 04 Oct 2023 16:29:38 GMT
content-type: text/html; charset=utf-8
set-cookie: 1e2670d92730b515ce3a1bb65da45062=1bf5e9573c9a2760c964ed1659cc1673; path=/; HttpOnly; Secure; SameSite=None
cache-control: private
```

3. Configure the DNS records for your cluster to target the front-end IP addresses of the user-managed load balancer. You must update records to your DNS server for the cluster API and applications over the load balancer.

Examples of modified DNS records

```
<load_balancer_ip_address> A api.<cluster_name>.<base_domain>
A record pointing to Load Balancer Front End
```

```
<load_balancer_ip_address> A apps.<cluster_name>.<base_domain>
A record pointing to Load Balancer Front End
```



IMPORTANT

DNS propagation might take some time for each DNS record to become available. Ensure that each DNS record propagates before validating each record.

4. For your OpenShift Container Platform cluster to use the user-managed load balancer, you must specify the following configuration in your cluster's **install-config.yaml** file:

```
# ...
platform:
  vsphere:
    loadBalancer:
      type: UserManaged 1
      apiVIPs:
        - <api_ip> 2
```

```
ingressVIPs:
- <ingress_ip> 3
# ...
```

- 1** Set **UserManaged** for the **type** parameter to specify a user-managed load balancer for your cluster. The parameter defaults to **OpenShiftManagedDefault**, which denotes the default internal load balancer. For services defined in an **openshift-kni-infra** namespace, a user-managed load balancer can deploy the **coredns** service to pods in your cluster but ignores **keepalived** and **haproxy** services.
- 2** Required parameter when you specify a user-managed load balancer. Specify the user-managed load balancer's public IP address, so that the Kubernetes API can communicate with the user-managed load balancer.
- 3** Required parameter when you specify a user-managed load balancer. Specify the user-managed load balancer's public IP address, so that the user-managed load balancer can manage ingress traffic for your cluster.

Verification

1. Use the **curl** CLI command to verify that the user-managed load balancer and DNS record configuration are operational:
 - a. Verify that you can access the cluster API, by running the following command and observing the output:

```
$ curl https://api.<cluster_name>.<base_domain>:6443/version --insecure
```

If the configuration is correct, you receive a JSON object in response:

```
{
  "major": "1",
  "minor": "11+",
  "gitVersion": "v1.11.0+ad103ed",
  "gitCommit": "ad103ed",
  "gitTreeState": "clean",
  "buildDate": "2019-01-09T06:44:10Z",
  "goVersion": "go1.10.3",
  "compiler": "gc",
  "platform": "linux/amd64"
}
```

- b. Verify that you can access the cluster machine configuration, by running the following command and observing the output:

```
$ curl -v https://api.<cluster_name>.<base_domain>:22623/healthz --insecure
```

If the configuration is correct, the output from the command shows the following response:

```
HTTP/1.1 200 OK
Content-Length: 0
```

- c. Verify that you can access each cluster application on port, by running the following command and observing the output:

```
$ curl http://console-openshift-console.apps.<cluster_name>.<base_domain> -I -L --insecure
```

If the configuration is correct, the output from the command shows the following response:

```
HTTP/1.1 302 Found
content-length: 0
location: https://console-openshift-console.apps.<cluster-name>.<base domain>/
cache-control: no-cacheHTTP/1.1 200 OK
referrer-policy: strict-origin-when-cross-origin
set-cookie: csrf-
token=39HoZgztDnzjJkq/JuLJMeoKNXIfiVv2YgZc09c3TBOBU4NI6kDXaJH1LdicNhN1UsQWzon4Dor9GWGfopaTEQ==; Path=/; Secure
x-content-type-options: nosniff
x-dns-prefetch-control: off
x-frame-options: DENY
x-xss-protection: 1; mode=block
date: Tue, 17 Nov 2020 08:42:10 GMT
content-type: text/html; charset=utf-8
set-cookie:
1e2670d92730b515ce3a1bb65da45062=9b714eb87e93cf34853e87a92d6894be; path=/; HttpOnly; Secure; SameSite=None
cache-control: private
```

- d. Verify that you can access each cluster application on port 443, by running the following command and observing the output:

```
$ curl https://console-openshift-console.apps.<cluster_name>.<base_domain> -I -L --insecure
```

If the configuration is correct, the output from the command shows the following response:

```
HTTP/1.1 200 OK
referrer-policy: strict-origin-when-cross-origin
set-cookie: csrf-
token=UIYWOyQ62LWjw2h003xtYSKlh1a0Py2hhctw0WmV2YEdhJfYqWwCGBsja261dGLgaYO0nxzVERhiXt6QepA7g==; Path=/; Secure; SameSite=Lax
x-content-type-options: nosniff
x-dns-prefetch-control: off
x-frame-options: DENY
x-xss-protection: 1; mode=block
date: Wed, 04 Oct 2023 16:29:38 GMT
content-type: text/html; charset=utf-8
set-cookie:
1e2670d92730b515ce3a1bb65da45062=1bf5e9573c9a2760c964ed1659cc1673; path=/; HttpOnly; Secure; SameSite=None
cache-control: private
```

23.2.5.9. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.



IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.
- Optional: Before you create the cluster, configure an external load balancer in place of the default load balancer.



IMPORTANT

You do not need to specify API and Ingress static addresses for your installation program. If you choose this configuration, you must take additional actions to define network targets that accept an IP address from each referenced vSphere subnet. See the section "Configuring a user-managed load balancer".

Procedure

- Change to the directory that contains the installation program and initialize the cluster deployment:


```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

- 1 For **<installation_directory>**, specify the location of your customized **./install-config.yaml** file.
- 2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to **<installation_directory>/openshift_install.log**.



IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.


Example output

...

```

INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s

```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

23.2.5.10. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```


23.2.5.11. Creating registry storage

After you install the cluster, you must create storage for the registry Operator.

23.2.5.11.1. Image registry removed during installation

On platforms that do not provide shareable object storage, the OpenShift Image Registry Operator bootstraps itself as **Removed**. This allows **openshift-installer** to complete installations on these platform types.

After installation, you must edit the Image Registry Operator configuration to switch the **managementState** from **Removed** to **Managed**. When this has completed, you must configure storage.

23.2.5.11.2. Image registry storage configuration

The Image Registry Operator is not initially available for platforms that do not provide default storage. After installation, you must configure your registry to use storage so that the Registry Operator is made available.

Instructions are shown for configuring a persistent volume, which is required for production clusters. Where applicable, instructions are shown for configuring an empty directory as the storage location, which is available for only non-production clusters.

Additional instructions are provided for allowing the image registry to use block storage types by using the **Recreate** rollout strategy during upgrades.

23.2.5.11.2.1. Configuring registry storage for VMware vSphere

As a cluster administrator, following installation you must configure your registry to use storage.

Prerequisites

- Cluster administrator permissions.
- A cluster on VMware vSphere.
- Persistent storage provisioned for your cluster, such as Red Hat OpenShift Data Foundation.



IMPORTANT

OpenShift Container Platform supports **ReadWriteOnce** access for image registry storage when you have only one replica. **ReadWriteOnce** access also requires that the registry uses the **Recreate** rollout strategy. To deploy an image registry that supports high availability with two or more replicas, **ReadWriteMany** access is required.

- Must have "100Gi" capacity.



IMPORTANT

Testing shows issues with using the NFS server on RHEL as storage backend for core services. This includes the OpenShift Container Registry and Quay, Prometheus for monitoring storage, and Elasticsearch for logging storage. Therefore, using RHEL NFS to back PVs used by core services is not recommended.

Other NFS implementations on the marketplace might not have these issues. Contact the individual NFS implementation vendor for more information on any testing that was possibly completed against these OpenShift Container Platform core components.

Procedure

1. To configure your registry to use storage, change the **spec.storage.pvc** in the **configs.imageregistry/cluster** resource.



NOTE


When you use shared storage, review your security settings to prevent outside access.

2. Verify that you do not have a registry pod:

```
$ oc get pod -n openshift-image-registry -l docker-registry=default
```

Example output

```
No resources found in openshift-image-registry namespace
```



NOTE

If you do have a registry pod in your output, you do not need to continue with this procedure.

3. Check the registry configuration:

```
$ oc edit configs.imageregistry.operator.openshift.io
```

Example output

```
storage:
  pvc:
    claim: 1
```

- 1** Leave the **claim** field blank to allow the automatic creation of an **image-registry-storage** persistent volume claim (PVC). The PVC is generated based on the default storage class. However, be aware that the default storage class might provide ReadWriteOnce (RWO) volumes, such as a RADOS Block Device (RBD), which can cause issues when you replicate to more than one replica.

4. Check the **clusteroperator** status:

-


```
$ oc get clusteroperator image-registry
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED
image-registry	4.7	True	False	6h50m

23.2.5.11.2.2. Configuring block registry storage for VMware vSphere

To allow the image registry to use block storage types such as vSphere Virtual Machine Disk (VMDK) during upgrades as a cluster administrator, you can use the **Recreate** rollout strategy.



IMPORTANT

Block storage volumes are supported but not recommended for use with image registry on production clusters. An installation where the registry is configured on block storage is not highly available because the registry cannot have more than one replica.

Procedure

1. Enter the following command to set the image registry storage as a block storage type, patch the registry so that it uses the **Recreate** rollout strategy, and runs with only **1** replica:

```
$ oc patch config.imageregistry.operator.openshift.io/cluster --type=merge -p '{"spec": {"rolloutStrategy": "Recreate", "replicas": 1}}'
```

2. Provision the PV for the block storage device, and create a PVC for that volume. The requested block volume uses the ReadWriteOnce (RWO) access mode.
 - a. Create a **pvc.yaml** file with the following contents to define a VMware vSphere **PersistentVolumeClaim** object:

```
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: image-registry-storage 1
  namespace: openshift-image-registry 2
spec:
  accessModes:
  - ReadWriteOnce 3
  resources:
    requests:
      storage: 100Gi 4
```

- 1 A unique name that represents the **PersistentVolumeClaim** object.
- 2 The namespace for the **PersistentVolumeClaim** object, which is **openshift-image-registry**.
- 3 The access mode of the persistent volume claim. With **ReadWriteOnce**, the volume can be mounted with read and write permissions by a single node.

4 The size of the persistent volume claim.

b. Enter the following command to create the **PersistentVolumeClaim** object from the file:

```
$ oc create -f pvc.yaml -n openshift-image-registry
```

3. Enter the following command to edit the registry configuration so that it references the correct PVC:

```
$ oc edit config.imageregistry.operator.openshift.io -o yaml
```

Example output

```
storage:
  pvc:
    claim: 1
```

1 By creating a custom PVC, you can leave the **claim** field blank for the default automatic creation of an **image-registry-storage** PVC.

For instructions about configuring registry storage so that it references the correct PVC, see [Configuring the registry for vSphere](#).

23.2.5.12. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service

23.2.5.13. Configuring network components to run on the control plane

You can configure networking components to run exclusively on the control plane nodes. By default, OpenShift Container Platform allows any node in the machine config pool to host the **ingressVIP** virtual IP address. However, some environments deploy compute nodes in separate subnets from the control plane nodes, which requires configuring the **ingressVIP** virtual IP address to run on the control plane nodes.



NOTE

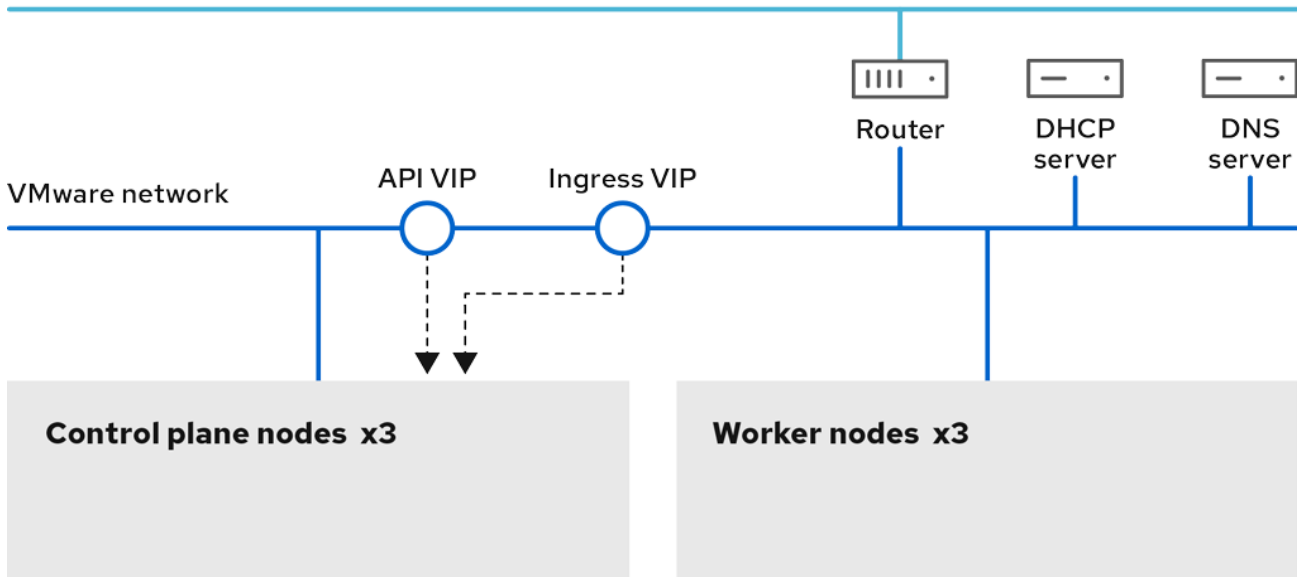
You can scale the remote nodes by creating a compute machine set in a separate subnet.



IMPORTANT

When deploying remote nodes in separate subnets, you must place the **ingressVIP** virtual IP address exclusively with the control plane nodes.

Internet access



Procedure

1. Change to the directory storing the **install-config.yaml** file:

```
$ cd ~/clusterconfigs
```

2. Switch to the **manifests** subdirectory:

```
$ cd manifests
```

3. Create a file named **cluster-network-avoid-workers-99-config.yaml**:

```
$ touch cluster-network-avoid-workers-99-config.yaml
```

4. Open the **cluster-network-avoid-workers-99-config.yaml** file in an editor and enter a custom resource (CR) that describes the Operator configuration:

```
apiVersion: machineconfiguration.openshift.io/v1
kind: MachineConfig
metadata:
  name: 50-worker-fix-ipi-rwn
  labels:
    machineconfiguration.openshift.io/role: worker
spec:
  config:
    ignition:
      version: 3.2.0
    storage:
```

```
files:
  - path: /etc/kubernetes/manifests/keepalived.yaml
    mode: 0644
    contents:
      source: data;
```

This manifest places the **ingressVIP** virtual IP address on the control plane nodes. Additionally, this manifest deploys the following processes on the control plane nodes only:

- **openshift-ingress-operator**
- **keepalived**

5. Save the **cluster-network-avoid-workers-99-config.yaml** file.
6. Create a **manifests/cluster-ingress-default-ingresscontroller.yaml** file:

```
apiVersion: operator.openshift.io/v1
kind: IngressController
metadata:
  name: default
  namespace: openshift-ingress-operator
spec:
  nodePlacement:
    nodeSelector:
      matchLabels:
        node-role.kubernetes.io/master: ""
```

7. Consider backing up the **manifests** directory. The installer deletes the **manifests/** directory when creating the cluster.
8. Modify the **cluster-scheduler-02-config.yml** manifest to make the control plane nodes schedulable by setting the **mastersSchedulable** field to **true**. Control plane nodes are not schedulable by default. For example:

```
$ sed -i "s;mastersSchedulable: false;mastersSchedulable: true;g"
clusterconfigs/manifests/cluster-scheduler-02-config.yml
```



NOTE

If control plane nodes are not schedulable after completing this procedure, deploying the cluster will fail.

23.2.5.14. Next steps

- [Customize your cluster.](#)
- If necessary, you can [opt out of remote health reporting](#) .
- [Set up your registry and configure registry storage](#) .
- Optional: [View the events from the vSphere Problem Detector Operator](#) to determine if the cluster has permission or storage configuration issues.

23.2.6. Installing a cluster on vSphere in a restricted network

In OpenShift Container Platform 4.16, you can install a cluster on VMware vSphere infrastructure in a restricted network by creating an internal mirror of the installation release content.



NOTE

OpenShift Container Platform supports deploying a cluster to a single VMware vCenter only. Deploying a cluster with machines/machine sets on multiple vCenters is not supported.

23.2.6.1. Prerequisites

- You have completed the tasks in [Preparing to install a cluster using installer-provisioned infrastructure](#).
- You reviewed your VMware platform licenses. Red Hat does not place any restrictions on your VMware licenses, but some VMware infrastructure components require licensing.
- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [created a registry on your mirror host](#) and obtained the **imageContentSources** data for your version of OpenShift Container Platform.



IMPORTANT

Because the installation media is on the mirror host, you can use that computer to complete all installation steps.

- You provisioned [persistent storage](#) for your cluster. To deploy a private image registry, your storage must provide the ReadWriteMany access mode.
- The OpenShift Container Platform installer requires access to port 443 on the vCenter and ESXi hosts. You verified that port 443 is accessible.
- If you use a firewall, you confirmed with the administrator that port 443 is accessible. Control plane nodes must be able to reach vCenter and ESXi hosts on port 443 for the installation to succeed.
- If you use a firewall and plan to use the Telemetry service, you [configured the firewall to allow the sites](#) that your cluster requires access to.



NOTE

If you are configuring a proxy, be sure to also review this site list.

23.2.6.2. About installations in restricted networks

In OpenShift Container Platform 4.16, you can perform an installation that does not require an active connection to the internet to obtain software components. Restricted network installations can be completed using installer-provisioned infrastructure or user-provisioned infrastructure, depending on

the cloud platform to which you are installing the cluster.

If you choose to perform a restricted network installation on a cloud platform, you still require access to its cloud APIs. Some cloud functions, like Amazon Web Service's Route 53 DNS and IAM services, require internet access. Depending on your network, you might require less internet access for an installation on bare metal hardware, Nutanix, or on VMware vSphere.

To complete a restricted network installation, you must create a registry that mirrors the contents of the OpenShift image registry and contains the installation media. You can create this registry on a mirror host, which can access both the internet and your closed network, or by using other methods that meet your restrictions.

23.2.6.2.1. Additional limits

Clusters in restricted networks have the following additional limitations and restrictions:

- The **ClusterVersion** status includes an **Unable to retrieve available updates** error.
- By default, you cannot use the contents of the Developer Catalog because you cannot access the required image stream tags.

23.2.6.3. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to obtain the images that are necessary to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.

23.2.6.4. Creating the RHCOS image for restricted network installations

Download the Red Hat Enterprise Linux CoreOS (RHCOS) image to install OpenShift Container Platform on a restricted network VMware vSphere environment.

Prerequisites

- Obtain the OpenShift Container Platform installation program. For a restricted network installation, the program is on your mirror registry host.

Procedure

1. Log in to the Red Hat Customer Portal's [Product Downloads page](#).
2. Under **Version**, select the most recent release of OpenShift Container Platform 4.16 for RHEL 8.



IMPORTANT


The RHCOS images might not change with every release of OpenShift Container Platform. You must download images with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Use the image versions that match your OpenShift Container Platform version if they are available.

3. Download the **Red Hat Enterprise Linux CoreOS (RHCOS) - vSphere** image.
4. Upload the image you downloaded to a location that is accessible from the bastion server.

The image is now available for a restricted installation. Note the image name or location for use in OpenShift Container Platform deployment.

23.2.6.5. VMware vSphere region and zone enablement

You can deploy an OpenShift Container Platform cluster to multiple vSphere datacenters that run in a single VMware vCenter. Each datacenter can run multiple clusters. This configuration reduces the risk of a hardware failure or network outage that can cause your cluster to fail. To enable regions and zones, you must define multiple failure domains for your OpenShift Container Platform cluster.



IMPORTANT

The VMware vSphere region and zone enablement feature requires the vSphere Container Storage Interface (CSI) driver as the default storage driver in the cluster. As a result, the feature is only available on a newly installed cluster.

For a cluster that was upgraded from a previous release, you must enable CSI automatic migration for the cluster. You can then configure multiple regions and zones for the upgraded cluster.

The default installation configuration deploys a cluster to a single vSphere datacenter. If you want to deploy a cluster to multiple vSphere datacenters, you must create an installation configuration file that enables the region and zone feature.

The default **install-config.yaml** file includes **vcenters** and **failureDomains** fields, where you can specify multiple vSphere datacenters and clusters for your OpenShift Container Platform cluster. You can leave these fields blank if you want to install an OpenShift Container Platform cluster in a vSphere environment that consists of single datacenter.

The following list describes terms associated with defining zones and regions for your cluster:

- **Failure domain:** Establishes the relationships between a region and zone. You define a failure domain by using vCenter objects, such as a **datastore** object. A failure domain defines the vCenter location for OpenShift Container Platform cluster nodes.
- **Region:** Specifies a vCenter datacenter. You define a region by using a tag from the **openshift-region** tag category.
- **Zone:** Specifies a vCenter cluster. You define a zone by using a tag from the **openshift-zone** tag category.

**NOTE**

If you plan on specifying more than one failure domain in your **install-config.yaml** file, you must create tag categories, zone tags, and region tags in advance of creating the configuration file.

You must create a vCenter tag for each vCenter datacenter, which represents a region. Additionally, you must create a vCenter tag for each cluster than runs in a datacenter, which represents a zone. After you create the tags, you must attach each tag to their respective datacenters and clusters.

The following table outlines an example of the relationship among regions, zones, and tags for a configuration with multiple vSphere datacenters running in a single VMware vCenter.

Datacenter (region)	Cluster (zone)	Tags
us-east	us-east-1	us-east-1a
		us-east-1b
	us-east-2	us-east-2a
		us-east-2b
us-west	us-west-1	us-west-1a
		us-west-1b
	us-west-2	us-west-2a
		us-west-2b

Additional resources

- [Additional VMware vSphere configuration parameters](#)
- [Deprecated VMware vSphere configuration parameters](#)
- [vSphere automatic migration](#)
- [VMware vSphere CSI Driver Operator](#)

23.2.6.6. Creating the installation configuration file

You can customize the OpenShift Container Platform cluster you install on VMware vSphere.

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster. For a restricted network installation, these files are on your mirror host.

- You have the **imageContentSources** values that were generated during mirror registry creation.
- You have obtained the contents of the certificate for your mirror registry.
- You have retrieved a Red Hat Enterprise Linux CoreOS (RHCOS) image and uploaded it to an accessible location.

Procedure

1. Create the **install-config.yaml** file.
 - a. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install create install-config --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the directory name to store the files that the installation program creates.

When specifying the directory:

- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
- Use an empty directory. Some installation assets, such as bootstrap X.509 certificates, have short expiration intervals, therefore you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

- b. At the prompts, provide the configuration details for your cloud:
 - i. Optional: Select an SSH key to use to access your cluster machines.



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- ii. Select **vsphere** as the platform to target.
- iii. Specify the name of your vCenter instance.
- iv. Specify the user name and password for the vCenter account that has the required permissions to create the cluster.
The installation program connects to your vCenter instance.
- v. Select the data center in your vCenter instance to connect to.

**NOTE**

After you create the installation configuration file, you can modify the file to create a multiple vSphere datacenters environment. This means that you can deploy an OpenShift Container Platform cluster to multiple vSphere datacenters that run in a single VMware vCenter. For more information about creating this environment, see the section named *VMware vSphere region and zone enablement*.

- vi. Select the default vCenter datastore to use.

**WARNING**

You can specify the path of any datastore that exists in a datastore cluster. By default, Storage Distributed Resource Scheduler (SDRS), which uses Storage vMotion, is automatically enabled for a datastore cluster. Red Hat does not support Storage vMotion, so you must disable Storage DRS to avoid data loss issues for your OpenShift Container Platform cluster.

You cannot specify more than one datastore path. If you must specify VMs across multiple datastores, use a **datastore** object to specify a failure domain in your cluster's **install-config.yaml** configuration file. For more information, see "VMware vSphere region and zone enablement".

- vii. Select the vCenter cluster to install the OpenShift Container Platform cluster in. The installation program uses the root resource pool of the vSphere cluster as the default resource pool.
 - viii. Select the network in the vCenter instance that contains the virtual IP addresses and DNS records that you configured.
 - ix. Enter the virtual IP address that you configured for control plane API access.
 - x. Enter the virtual IP address that you configured for cluster ingress.
 - xi. Enter the base domain. This base domain must be the same one that you used in the DNS records that you configured.
 - xii. Enter a descriptive name for your cluster.
The cluster name you enter must match the cluster name you specified when configuring the DNS records.
2. In the **install-config.yaml** file, set the value of **platform.vsphere.clusterOSImage** to the image location or name. For example:

```
platform:
  vsphere:
    clusterOSImage: http://mirror.example.com/images/rhcos-43.81.201912131630.0-
```

```
vmware.x86_64.ova?
sha256=ffebbd68e8a1f2a245ca19522c16c86f67f9ac8e4e0c1f0a812b068b16f7265d
```

3. Edit the **install-config.yaml** file to give the additional information that is required for an installation in a restricted network.

- a. Update the **pullSecret** value to contain the authentication information for your registry:

```
pullSecret: '{"auths":{"<mirror_host_name>:5000": {"auth": "<credentials>","email":
"you@example.com"}}}'
```

For **<mirror_host_name>**, specify the registry domain name that you specified in the certificate for your mirror registry, and for **<credentials>**, specify the base64-encoded user name and password for your mirror registry.

- b. Add the **additionalTrustBundle** parameter and value.

```
additionalTrustBundle: |
  -----BEGIN CERTIFICATE-----

  /-----END CERTIFICATE-----
```

The value must be the contents of the certificate file that you used for your mirror registry. The certificate file can be an existing, trusted certificate authority, or the self-signed certificate that you generated for the mirror registry.

- c. Add the image content resources, which resemble the following YAML excerpt:

```
imageContentSources:
- mirrors:
  - <mirror_host_name>:5000/<repo_name>/release
  source: quay.io/openshift-release-dev/ocp-release
- mirrors:
  - <mirror_host_name>:5000/<repo_name>/release
  source: registry.redhat.io/ocp/release
```


For these values, use the **imageContentSources** that you recorded during mirror registry creation.

- d. Optional: Set the publishing strategy to **Internal**:

```
publish: Internal
```

By setting this option, you create an internal Ingress Controller and a private load balancer.

4. Make any other modifications to the **install-config.yaml** file that you require. For more information about the parameters, see "Installation configuration parameters".
5. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

Additional resources

- [Installation configuration parameters](#)

23.2.6.6.1. Sample install-config.yaml file for an installer-provisioned VMware vSphere cluster

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.

```

apiVersion: v1
baseDomain: example.com 1
compute: 2
- architecture: amd64
  name: <worker_node>
  platform: {}
  replicas: 3
controlPlane: 3
  architecture: amd64
  name: <parent_node>
  platform: {}
  replicas: 3
metadata:
  creationTimestamp: null
  name: test 4
platform:
  vsphere: 5
    apiVIPs:
      - 10.0.0.1
    failureDomains: 6
      - name: <failure_domain_name>
        region: <default_region_name>
        server: <fully_qualified_domain_name>
    topology:
      computeCluster: "/<datacenter>/host/<cluster>"
      datacenter: <datacenter>
      datastore: "/<datacenter>/datastore/<datastore>" 7
      networks:
        - <VM_Network_name>
      resourcePool: "/<datacenter>/host/<cluster>/Resources/<resourcePool>" 8
      folder: "/<datacenter_name>/vm/<folder_name>/<subfolder_name>"
      tagIDs: 9
        - <tag_id> 10
      zone: <default_zone_name>
    ingressVIPs:
      - 10.0.0.2
  vcenters:
    - datacenters:
      - <datacenter>
      password: <password>
      port: 443
      server: <fully_qualified_domain_name>
      user: administrator@vsphere.local
    diskType: thin 11
    clusterOSImage: http://mirror.example.com/images/rhcos-47.83.202103221318-0-

```


The ID of the tag to be associated by the installation program. For example, **urn:vmomi:InventoryServiceTag:208e713c-cae3-4b7f-918e-4051ca7d1f97:GLOBAL**. For more

- 11 The vSphere disk provisioning method.
- 12 The location of the Red Hat Enterprise Linux CoreOS (RHCOS) image that is accessible from the bastion server.
- 13 For **<local_registry>**, specify the registry domain name, and optionally the port, that your mirror registry uses to serve content. For example **registry.example.com** or **registry.example.com:5000**. For **<credentials>**, specify the base64-encoded user name and password for your mirror registry.
- 14 Provide the contents of the certificate file that you used for your mirror registry.
- 15 Provide the **imageContentSources** section from the output of the command to mirror the repository.



NOTE


In OpenShift Container Platform 4.12 and later, the **apiVIP** and **ingressVIP** configuration settings are deprecated. Instead, use a list format to enter values in the **apiVIPs** and **ingressVIPs** configuration settings.

23.2.6.6.2. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure


1. Edit your **install-config.yaml** file and add the proxy settings. For example:


```

apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
additionalTrustBundle: | 4
  -----BEGIN CERTIFICATE-----
  <MY_TRUSTED_CA_CERT>
  -----END CERTIFICATE-----
additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5

```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations. You must include vCenter's IP address and the IP range that you use for its machines.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.

**NOTE**

The installation program does not support the proxy **readinessEndpoints** field.

**NOTE**

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.


**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

23.2.6.6.3. Configuring regions and zones for a VMware vCenter

You can modify the default installation configuration file, so that you can deploy an OpenShift Container Platform cluster to multiple vSphere datacenters that run in a single VMware vCenter.

The default **install-config.yaml** file configuration from the previous release of OpenShift Container Platform is deprecated. You can continue to use the deprecated default configuration, but the **openshift-installer** will prompt you with a warning message that indicates the use of deprecated fields in the configuration file.

**IMPORTANT**

The example uses the **govc** command. The **govc** command is an open source command available from VMware; it is not available from Red Hat. The Red Hat support team does not maintain the **govc** command. Instructions for downloading and installing **govc** are found on the VMware documentation website

Prerequisites

- You have an existing **install-config.yaml** installation configuration file.

**IMPORTANT**

You must specify at least one failure domain for your OpenShift Container Platform cluster, so that you can provision datacenter objects for your VMware vCenter server. Consider specifying multiple failure domains if you need to provision virtual machine nodes in different datacenters, clusters, datastores, and other components. To enable regions and zones, you must define multiple failure domains for your OpenShift Container Platform cluster.

Procedure

1. Enter the following **govc** command-line tool commands to create the **openshift-region** and **openshift-zone** vCenter tag categories:

**IMPORTANT**

If you specify different names for the **openshift-region** and **openshift-zone** vCenter tag categories, the installation of the OpenShift Container Platform cluster fails.

```
$ govc tags.category.create -d "OpenShift region" openshift-region
```

```
$ govc tags.category.create -d "OpenShift zone" openshift-zone
```

2. To create a region tag for each region vSphere datacenter where you want to deploy your cluster, enter the following command in your terminal:

```
$ govc tags.create -c <region_tag_category> <region_tag>
```

- To create a zone tag for each vSphere cluster where you want to deploy your cluster, enter the following command:

```
$ govc tags.create -c <zone_tag_category> <zone_tag>
```

- Attach region tags to each vCenter datacenter object by entering the following command:

```
$ govc tags.attach -c <region_tag_category> <region_tag_1> /<datacenter_1>
```

- Attach the zone tags to each vCenter datacenter object by entering the following command:

```
$ govc tags.attach -c <zone_tag_category> <zone_tag_1> /<datacenter_1>/host/vcs-mdcnc-workload-1
```

- Change to the directory that contains the installation program and initialize the cluster deployment according to your chosen installation requirements.

Sample install-config.yaml file with multiple datacenters defined in a vSphere center

```
---
compute:
---
vsphere:
  zones:
    - "<machine_pool_zone_1>"
    - "<machine_pool_zone_2>"
---
controlPlane:
---
vsphere:
  zones:
    - "<machine_pool_zone_1>"
    - "<machine_pool_zone_2>"
---
platform:
  vsphere:
    vcenters:
---
  datacenters:
    - <datacenter1_name>
    - <datacenter2_name>
  failureDomains:
    - name: <machine_pool_zone_1>
      region: <region_tag_1>
      zone: <zone_tag_1>
      server: <fully_qualified_domain_name>
  topology:
    datacenter: <datacenter1>
    computeCluster: "/<datacenter1>/host/<cluster1>"
    networks:
      - <VM_Network1_name>
    datastore: "/<datacenter1>/datastore/<datastore1>"
```

```

resourcePool: "/<datacenter1>/host/<cluster1>/Resources/<resourcePool1>"
folder: "/<datacenter1>/vm/<folder1>"
- name: <machine_pool_zone_2>
  region: <region_tag_2>
  zone: <zone_tag_2>
  server: <fully_qualified_domain_name>
  topology:
    datacenter: <datacenter2>
    computeCluster: "/<datacenter2>/host/<cluster2>"
    networks:
      - <VM_Network2_name>
    datastore: "/<datacenter2>/datastore/<datastore2>"
    resourcePool: "/<datacenter2>/host/<cluster2>/Resources/<resourcePool2>"
    folder: "/<datacenter2>/vm/<folder2>"
---
```

23.2.6.7. Services for a user-managed load balancer

You can configure an OpenShift Container Platform cluster to use a user-managed load balancer in place of the default load balancer.



IMPORTANT

Configuring a user-managed load balancer depends on your vendor's load balancer.

The information and examples in this section are for guideline purposes only. Consult the vendor documentation for more specific information about the vendor's load balancer.

Red Hat supports the following services for a user-managed load balancer:

- Ingress Controller
- OpenShift API
- OpenShift MachineConfig API

You can choose whether you want to configure one or all of these services for a user-managed load balancer. Configuring only the Ingress Controller service is a common configuration option. To better understand each service, view the following diagrams:

Figure 23.7. Example network workflow that shows an Ingress Controller operating in an OpenShift Container Platform environment




Figure 23.8. Example network workflow that shows an OpenShift API operating in an OpenShift Container Platform environment

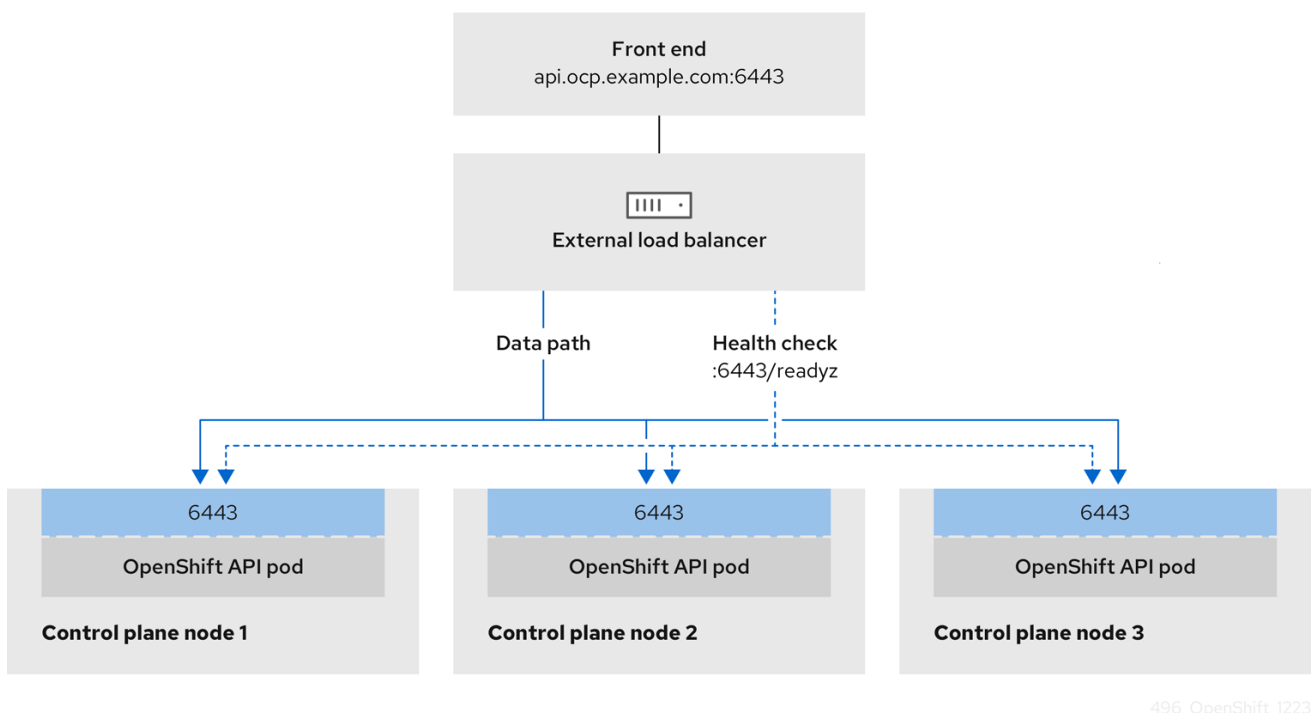
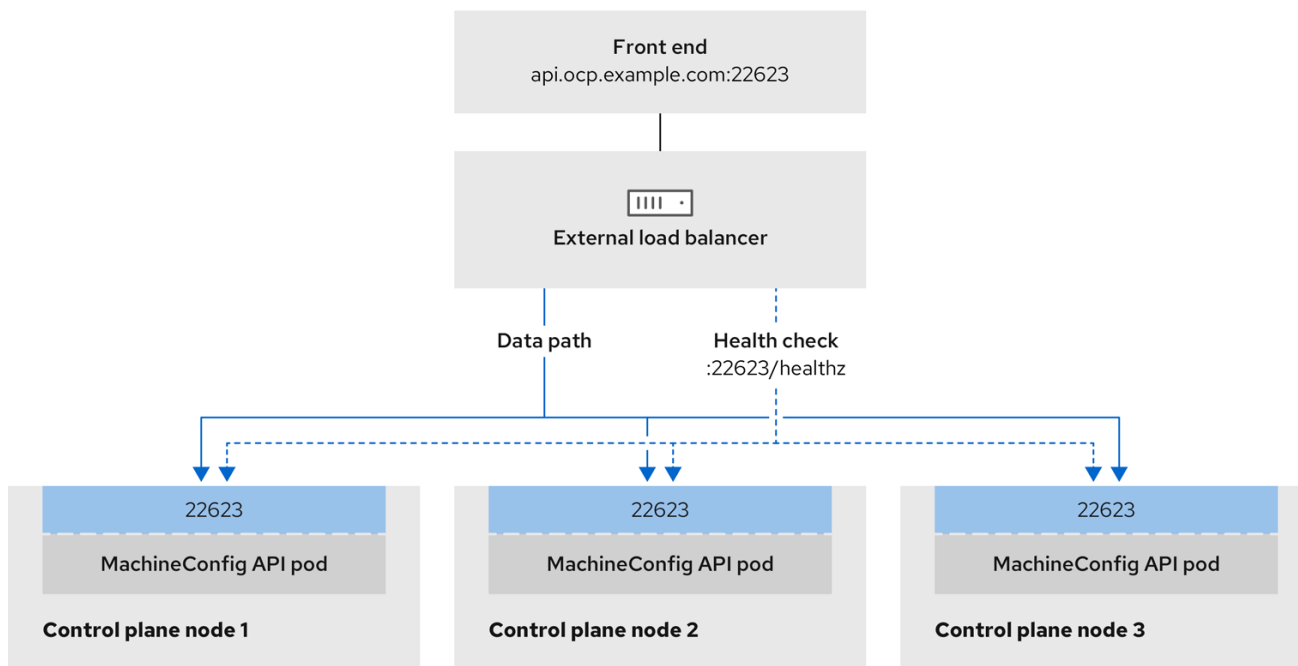


Figure 23.9. Example network workflow that shows an OpenShift MachineConfig API operating in an OpenShift Container Platform environment



496_OpenShift_1223

The following configuration options are supported for user-managed load balancers:

- Use a node selector to map the Ingress Controller to a specific set of nodes. You must assign a static IP address to each node in this set, or configure each node to receive the same IP address from the Dynamic Host Configuration Protocol (DHCP). Infrastructure nodes commonly receive this type of configuration.
- Target all IP addresses on a subnet. This configuration can reduce maintenance overhead, because you can create and destroy nodes within those networks without reconfiguring the load balancer targets. If you deploy your ingress pods by using a machine set on a smaller network, such as a `/27` or `/28`, you can simplify your load balancer targets.

TIP

You can list all IP addresses that exist in a network by checking the machine config pool's resources.


Before you configure a user-managed load balancer for your OpenShift Container Platform cluster, consider the following information:

- For a front-end IP address, you can use the same IP address for the front-end IP address, the Ingress Controller's load balancer, and API load balancer. Check the vendor's documentation for this capability.
- For a back-end IP address, ensure that an IP address for an OpenShift Container Platform control plane node does not change during the lifetime of the user-managed load balancer. You can achieve this by completing one of the following actions:
 - Assign a static IP address to each control plane node.

- Configure each node to receive the same IP address from the DHCP every time the node requests a DHCP lease. Depending on the vendor, the DHCP lease might be in the form of an IP reservation or a static DHCP assignment.
- Manually define each node that runs the Ingress Controller in the user-managed load balancer for the Ingress Controller back-end service. For example, if the Ingress Controller moves to an undefined node, a connection outage can occur.

23.2.6.7.1. Configuring a user-managed load balancer

You can configure an OpenShift Container Platform cluster to use a user-managed load balancer in place of the default load balancer.



IMPORTANT

Before you configure a user-managed load balancer, ensure that you read the "Services for a user-managed load balancer" section.

Read the following prerequisites that apply to the service that you want to configure for your user-managed load balancer.



NOTE

MetalLB, which runs on a cluster, functions as a user-managed load balancer.

OpenShift API prerequisites

- You defined a front-end IP address.
- TCP ports 6443 and 22623 are exposed on the front-end IP address of your load balancer. Check the following items:
 - Port 6443 provides access to the OpenShift API service.
 - Port 22623 can provide ignition startup configurations to nodes.
- The front-end IP address and port 6443 are reachable by all users of your system with a location external to your OpenShift Container Platform cluster.
- The front-end IP address and port 22623 are reachable only by OpenShift Container Platform nodes.
- The load balancer backend can communicate with OpenShift Container Platform control plane nodes on port 6443 and 22623.

Ingress Controller prerequisites

- You defined a front-end IP address.
- TCP ports 443 and 80 are exposed on the front-end IP address of your load balancer.
- The front-end IP address, port 80 and port 443 are be reachable by all users of your system with a location external to your OpenShift Container Platform cluster.

- The front-end IP address, port 80 and port 443 are reachable to all nodes that operate in your OpenShift Container Platform cluster.
- The load balancer backend can communicate with OpenShift Container Platform nodes that run the Ingress Controller on ports 80, 443, and 1936.

Prerequisite for health check URL specifications

You can configure most load balancers by setting health check URLs that determine if a service is available or unavailable. OpenShift Container Platform provides these health checks for the OpenShift API, Machine Configuration API, and Ingress Controller backend services.

The following examples show health check specifications for the previously listed backend services:

Example of a Kubernetes API health check specification

```
Path: HTTPS:6443/readyz
Healthy threshold: 2
Unhealthy threshold: 2
Timeout: 10
Interval: 10
```

Example of a Machine Config API health check specification

```
Path: HTTPS:22623/healthz
Healthy threshold: 2
Unhealthy threshold: 2
Timeout: 10
Interval: 10
```

Example of an Ingress Controller health check specification

```
Path: HTTP:1936/healthz/ready
Healthy threshold: 2
Unhealthy threshold: 2
Timeout: 5
Interval: 10
```

Procedure

1. Configure the HAProxy Ingress Controller, so that you can enable access to the cluster from your load balancer on ports 6443, 22623, 443, and 80. Depending on your needs, you can specify the IP address of a single subnet or IP addresses from multiple subnets in your HAProxy configuration.

Example HAProxy configuration with one listed subnet

```
# ...
listen my-cluster-api-6443
  bind 192.168.1.100:6443
  mode tcp
  balance roundrobin
  option httpchk
  http-check connect
  http-check send meth GET uri /readyz
```



```

http-check expect status 200
  server my-cluster-master-2 192.168.1.101:6443 check inter 10s rise 2 fall 2
  server my-cluster-master-0 192.168.1.102:6443 check inter 10s rise 2 fall 2
  server my-cluster-master-1 192.168.1.103:6443 check inter 10s rise 2 fall 2

listen my-cluster-machine-config-api-22623
  bind 192.168.1.100:22623
  mode tcp
  balance roundrobin
  option httpchk
  http-check connect
  http-check send meth GET uri /healthz
  http-check expect status 200
    server my-cluster-master-2 192.168.1.101:22623 check inter 10s rise 2 fall 2
    server my-cluster-master-0 192.168.1.102:22623 check inter 10s rise 2 fall 2
    server my-cluster-master-1 192.168.1.103:22623 check inter 10s rise 2 fall 2

listen my-cluster-apps-443
  bind 192.168.1.100:443
  mode tcp
  balance roundrobin
  option httpchk
  http-check connect
  http-check send meth GET uri /healthz/ready
  http-check expect status 200
    server my-cluster-worker-0 192.168.1.111:443 check port 1936 inter 10s rise 2 fall 2
    server my-cluster-worker-1 192.168.1.112:443 check port 1936 inter 10s rise 2 fall 2
    server my-cluster-worker-2 192.168.1.113:443 check port 1936 inter 10s rise 2 fall 2

listen my-cluster-apps-80
  bind 192.168.1.100:80
  mode tcp
  balance roundrobin
  option httpchk
  http-check connect
  http-check send meth GET uri /healthz/ready
  http-check expect status 200
    server my-cluster-worker-0 192.168.1.111:80 check port 1936 inter 10s rise 2 fall 2
    server my-cluster-worker-1 192.168.1.112:80 check port 1936 inter 10s rise 2 fall 2
    server my-cluster-worker-2 192.168.1.113:80 check port 1936 inter 10s rise 2 fall 2
# ...

```

Example HAProxy configuration with multiple listed subnets

```

# ...
listen api-server-6443
  bind *:6443
  mode tcp
    server master-00 192.168.83.89:6443 check inter 1s
    server master-01 192.168.84.90:6443 check inter 1s
    server master-02 192.168.85.99:6443 check inter 1s
    server bootstrap 192.168.80.89:6443 check inter 1s

listen machine-config-server-22623
  bind *:22623
  mode tcp

```

```

server master-00 192.168.83.89:22623 check inter 1s
server master-01 192.168.84.90:22623 check inter 1s
server master-02 192.168.85.99:22623 check inter 1s
server bootstrap 192.168.80.89:22623 check inter 1s

listen ingress-router-80
  bind *:80
  mode tcp
  balance source
    server worker-00 192.168.83.100:80 check inter 1s
    server worker-01 192.168.83.101:80 check inter 1s

listen ingress-router-443
  bind *:443
  mode tcp
  balance source
    server worker-00 192.168.83.100:443 check inter 1s
    server worker-01 192.168.83.101:443 check inter 1s

listen ironic-api-6385
  bind *:6385
  mode tcp
  balance source
    server master-00 192.168.83.89:6385 check inter 1s
    server master-01 192.168.84.90:6385 check inter 1s
    server master-02 192.168.85.99:6385 check inter 1s
    server bootstrap 192.168.80.89:6385 check inter 1s

listen inspector-api-5050
  bind *:5050
  mode tcp
  balance source
    server master-00 192.168.83.89:5050 check inter 1s
    server master-01 192.168.84.90:5050 check inter 1s
    server master-02 192.168.85.99:5050 check inter 1s
    server bootstrap 192.168.80.89:5050 check inter 1s
# ...

```

2. Use the **curl** CLI command to verify that the user-managed load balancer and its resources are operational:
 - a. Verify that the cluster machine configuration API is accessible to the Kubernetes API server resource, by running the following command and observing the response:

```
$ curl https://<loadbalancer_ip_address>:6443/version --insecure
```

If the configuration is correct, you receive a JSON object in response:

```
{
  "major": "1",
  "minor": "11+",
  "gitVersion": "v1.11.0+ad103ed",
  "gitCommit": "ad103ed",
  "gitTreeState": "clean",
  "buildDate": "2019-01-09T06:44:10Z",
  "goVersion": "go1.10.3",
}
```

```
"compiler": "gc",
"platform": "linux/amd64"
}
```

- b. Verify that the cluster machine configuration API is accessible to the Machine config server resource, by running the following command and observing the output:

```
$ curl -v https://<loadbalancer_ip_address>:22623/healthz --insecure
```

If the configuration is correct, the output from the command shows the following response:

```
HTTP/1.1 200 OK
Content-Length: 0
```

- c. Verify that the controller is accessible to the Ingress Controller resource on port 80, by running the following command and observing the output:

```
$ curl -I -L -H "Host: console-openshift-console.apps.<cluster_name>.<base_domain>"
http://<load_balancer_front_end_IP_address>
```

If the configuration is correct, the output from the command shows the following response:

```
HTTP/1.1 302 Found
content-length: 0
location: https://console-openshift-console.apps.ocp4.private.opequon.net/
cache-control: no-cache
```

- d. Verify that the controller is accessible to the Ingress Controller resource on port 443, by running the following command and observing the output:

```
$ curl -I -L --insecure --resolve console-openshift-console.apps.<cluster_name>.<base_domain>:443:<Load Balancer Front End IP Address> https://console-openshift-console.apps.<cluster_name>.<base_domain>
```

If the configuration is correct, the output from the command shows the following response:


```
HTTP/1.1 200 OK
referrer-policy: strict-origin-when-cross-origin
set-cookie: csrf-
token=UIYWOyQ62LWjw2h003xtYSKlh1a0Py2hhctw0WmV2YEdhJfYqWwCGBsja261dG
LgaYO0nxzVERhiXt6QepA7g==; Path=/; Secure; SameSite=Lax
x-content-type-options: nosniff
x-dns-prefetch-control: off
x-frame-options: DENY
x-xss-protection: 1; mode=block
date: Wed, 04 Oct 2023 16:29:38 GMT
content-type: text/html; charset=utf-8
set-cookie:
1e2670d92730b515ce3a1bb65da45062=1bf5e9573c9a2760c964ed1659cc1673; path=/;
HttpOnly; Secure; SameSite=None
cache-control: private
```

- Configure the DNS records for your cluster to target the front-end IP addresses of the user-managed load balancer. You must update records to your DNS server for the cluster API and applications over the load balancer.

Examples of modified DNS records

```
<load_balancer_ip_address> A api.<cluster_name>.<base_domain>
A record pointing to Load Balancer Front End
```

```
<load_balancer_ip_address> A apps.<cluster_name>.<base_domain>
A record pointing to Load Balancer Front End
```



IMPORTANT

DNS propagation might take some time for each DNS record to become available. Ensure that each DNS record propagates before validating each record.

- For your OpenShift Container Platform cluster to use the user-managed load balancer, you must specify the following configuration in your cluster's **install-config.yaml** file:

```
# ...
platform:
  vsphere:
    loadBalancer:
      type: UserManaged 1
      apiVIPs:
        - <api_ip> 2
      ingressVIPs:
        - <ingress_ip> 3
# ...
```

- Set **UserManaged** for the **type** parameter to specify a user-managed load balancer for your cluster. The parameter defaults to **OpenShiftManagedDefault**, which denotes the default internal load balancer. For services defined in an **openshift-kni-infra** namespace, a user-managed load balancer can deploy the **coredns** service to pods in your cluster but ignores **keepalived** and **haproxy** services.
- Required parameter when you specify a user-managed load balancer. Specify the user-managed load balancer's public IP address, so that the Kubernetes API can communicate with the user-managed load balancer.
- Required parameter when you specify a user-managed load balancer. Specify the user-managed load balancer's public IP address, so that the user-managed load balancer can manage ingress traffic for your cluster.

Verification

- Use the **curl** CLI command to verify that the user-managed load balancer and DNS record configuration are operational:
 - Verify that you can access the cluster API, by running the following command and observing the output:

```
$ curl https://api.<cluster_name>.<base_domain>:6443/version --insecure
```

If the configuration is correct, you receive a JSON object in response:

```
{
  "major": "1",
  "minor": "11+",
  "gitVersion": "v1.11.0+ad103ed",
  "gitCommit": "ad103ed",
  "gitTreeState": "clean",
  "buildDate": "2019-01-09T06:44:10Z",
  "goVersion": "go1.10.3",
  "compiler": "gc",
  "platform": "linux/amd64"
}
```

- b. Verify that you can access the cluster machine configuration, by running the following command and observing the output:

```
$ curl -v https://api.<cluster_name>.<base_domain>:22623/healthz --insecure
```

If the configuration is correct, the output from the command shows the following response:

```
HTTP/1.1 200 OK
Content-Length: 0
```

- c. Verify that you can access each cluster application on port, by running the following command and observing the output:

```
$ curl http://console-openshift-console.apps.<cluster_name>.<base_domain> -I -L --insecure
```

If the configuration is correct, the output from the command shows the following response:

```
HTTP/1.1 302 Found
content-length: 0
location: https://console-openshift-console.apps.<cluster-name>.<base domain>/
cache-control: no-cacheHTTP/1.1 200 OK
referrer-policy: strict-origin-when-cross-origin
set-cookie: csrf-
token=39HoZgztDnzjJkq/JuLJMeoKNXIfiVv2YgZc09c3TBOBU4NI6kDXaJH1LdicNhN1UsQ
Wzon4Dor9GWGfopaTEQ==; Path=/; Secure
x-content-type-options: nosniff
x-dns-prefetch-control: off
x-frame-options: DENY
x-xss-protection: 1; mode=block
date: Tue, 17 Nov 2020 08:42:10 GMT
content-type: text/html; charset=utf-8
set-cookie:
1e2670d92730b515ce3a1bb65da45062=9b714eb87e93cf34853e87a92d6894be; path=/;
HttpOnly; Secure; SameSite=None
cache-control: private
```

- d. Verify that you can access each cluster application on port 443, by running the following command and observing the output:


```
$ curl https://console-openshift-console.apps.<cluster_name>.<base_domain> -I -L --insecure
```

If the configuration is correct, the output from the command shows the following response:

```
HTTP/1.1 200 OK
referrer-policy: strict-origin-when-cross-origin
set-cookie: csrf-
token=UIYWOyQ62LWjw2h003xtYSKlh1a0Py2hhctw0WmV2YEdhJjFyQwWcGBsja261dG
LgaYO0nxzVERhiXt6QepA7g==; Path=/; Secure; SameSite=Lax
x-content-type-options: nosniff
x-dns-prefetch-control: off
x-frame-options: DENY
x-xss-protection: 1; mode=block
date: Wed, 04 Oct 2023 16:29:38 GMT
content-type: text/html; charset=utf-8
set-cookie:
1e2670d92730b515ce3a1bb65da45062=1bf5e9573c9a2760c964ed1659cc1673; path=/;
HttpOnly; Secure; SameSite=None
cache-control: private
```

23.2.6.8. Deploying the cluster

You can install OpenShift Container Platform on a compatible cloud platform.




IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

Prerequisites

- You have the OpenShift Container Platform installation program and the pull secret for your cluster.
- You have verified that the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.
- Optional: Before you create the cluster, configure an external load balancer in place of the default load balancer.



IMPORTANT

You do not need to specify API and Ingress static addresses for your installation program. If you choose this configuration, you must take additional actions to define network targets that accept an IP address from each referenced vSphere subnet. See the section "Configuring a user-managed load balancer".

Procedure

- Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

- 1 For **<installation_directory>**, specify the location of your customized **./install-config.yaml** file.
- 2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to **<installation_directory>/openshift_install.log**.



IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

23.2.6.9. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

23.2.6.10. Disabling the default OperatorHub catalog sources

Operator catalogs that source content provided by Red Hat and community projects are configured for OperatorHub by default during an OpenShift Container Platform installation. In a restricted network environment, you must disable the default catalogs as a cluster administrator.

Procedure

- Disable the sources for the default catalogs by adding **disableAllDefaultSources: true** to the **OperatorHub** object:

```
$ oc patch OperatorHub cluster --type json \  
-p '[{"op": "add", "path": "/spec/disableAllDefaultSources", "value": true}]'
```

TIP

Alternatively, you can use the web console to manage catalog sources. From the **Administration** → **Cluster Settings** → **Configuration** → **OperatorHub** page, click the **Sources** tab, where you can create, update, delete, disable, and enable individual sources.

23.2.6.11. Creating registry storage

After you install the cluster, you must create storage for the Registry Operator.

23.2.6.11.1. Image registry removed during installation

On platforms that do not provide shareable object storage, the OpenShift Image Registry Operator bootstraps itself as **Removed**. This allows **openshift-installer** to complete installations on these platform types.

After installation, you must edit the Image Registry Operator configuration to switch the **managementState** from **Removed** to **Managed**. When this has completed, you must configure storage.

23.2.6.11.2. Image registry storage configuration

The Image Registry Operator is not initially available for platforms that do not provide default storage. After installation, you must configure your registry to use storage so that the Registry Operator is made available.

Instructions are shown for configuring a persistent volume, which is required for production clusters. Where applicable, instructions are shown for configuring an empty directory as the storage location, which is available for only non-production clusters.


Additional instructions are provided for allowing the image registry to use block storage types by using the **Recreate** rollout strategy during upgrades.

23.2.6.11.2.1. Configuring registry storage for VMware vSphere

As a cluster administrator, following installation you must configure your registry to use storage.

Prerequisites

- Cluster administrator permissions.
- A cluster on VMware vSphere.
- Persistent storage provisioned for your cluster, such as Red Hat OpenShift Data Foundation.



IMPORTANT

OpenShift Container Platform supports **ReadWriteOnce** access for image registry storage when you have only one replica. **ReadWriteOnce** access also requires that the registry uses the **Recreate** rollout strategy. To deploy an image registry that supports high availability with two or more replicas, **ReadWriteMany** access is required.

- Must have "100Gi" capacity.



IMPORTANT

Testing shows issues with using the NFS server on RHEL as storage backend for core services. This includes the OpenShift Container Registry and Quay, Prometheus for monitoring storage, and Elasticsearch for logging storage. Therefore, using RHEL NFS to back PVs used by core services is not recommended.

Other NFS implementations on the marketplace might not have these issues. Contact the individual NFS implementation vendor for more information on any testing that was possibly completed against these OpenShift Container Platform core components.

Procedure

1. To configure your registry to use storage, change the **spec.storage.pvc** in the **configs.imageregistry/cluster** resource.



NOTE

When you use shared storage, review your security settings to prevent outside access.

2. Verify that you do not have a registry pod:

```
$ oc get pod -n openshift-image-registry -l docker-registry=default
```

Example output

```
No resources found in openshift-image-registry namespace
```



NOTE

If you do have a registry pod in your output, you do not need to continue with this procedure.

3. Check the registry configuration:

```
$ oc edit configs.imageregistry.operator.openshift.io
```

Example output

```
storage:
  pvc:
    claim: 1
```

- 1** Leave the **claim** field blank to allow the automatic creation of an **image-registry-storage** persistent volume claim (PVC). The PVC is generated based on the default storage class. However, be aware that the default storage class might provide ReadWriteOnce (RWO) volumes, such as a RADOS Block Device (RBD), which can cause issues when you replicate to more than one replica.

4. Check the **clusteroperator** status:

```
$ oc get clusteroperator image-registry
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED
image-registry	4.7	True	False	False

23.2.6.12. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service

23.2.6.13. Next steps

- [Customize your cluster](#).
- If necessary, you can [opt out of remote health reporting](#).
- If necessary, see [Registering your disconnected cluster](#).
- [Set up your registry and configure registry storage](#).

23.3. USER-PROVISIONED INFRASTRUCTURE

23.3.1. vSphere installation requirements for user-provisioned infrastructure

Before you begin an installation on infrastructure that you provision, be sure that your vSphere environment meets the following installation requirements.

23.3.1.1. VMware vSphere infrastructure requirements

You must install an OpenShift Container Platform cluster on one of the following versions of a VMware vSphere instance that meets the requirements for the components that you use:


- Version 7.0 Update 2 or later
- Version 8.0 Update 1 or later

Both of these releases support Container Storage Interface (CSI) migration, which is enabled by default on OpenShift Container Platform 4.16.

You can host the VMware vSphere infrastructure on-premise or on a [VMware Cloud Verified provider](#) that meets the requirements outlined in the following tables:

Table 23.18. Version requirements for vSphere virtual environments

Virtual environment product	Required version
VMware virtual hardware	15 or later
vSphere ESXi hosts	7.0 Update 2 or later; 8.0 Update 1 or later
vCenter host	7.0 Update 2 or later; 8.0 Update 1 or later

**IMPORTANT**

You must ensure that the time on your ESXi hosts is synchronized before you install OpenShift Container Platform. See [Edit Time Configuration for a Host](#) in the VMware documentation.

Table 23.19. Minimum supported vSphere version for VMware components

Component	Minimum supported versions	Description
Hypervisor	vSphere 7.0 Update 2 (or later) or vSphere 8.0 Update 1 (or later) with virtual hardware version 15	This hypervisor version is the minimum version that Red Hat Enterprise Linux CoreOS (RHCOS) supports. For more information about supported hardware on the latest version of Red Hat Enterprise Linux (RHEL) that is compatible with RHCOS, see Hardware on the Red Hat Customer Portal.
Optional: Networking (NSX-T)	vSphere 7.0 Update 2 or later; vSphere 8.0 Update 1 or later	At a minimum, vSphere 7.0 Update 2 or vSphere 8.0 Update 1 is required for OpenShift Container Platform. For more information about the compatibility of NSX and OpenShift Container Platform, see the Release Notes section of VMware's NSX container plugin documentation .
CPU micro-architecture	x86-64-v2 or higher	OpenShift 4.13 and later are based on RHEL 9.2 host operating system which raised the microarchitecture requirements to x86-64-v2. See the RHEL Microarchitecture requirements documentation . You can verify compatibility by following the procedures outlined in this KCS article .



IMPORTANT

To ensure the best performance conditions for your cluster workloads that operate on Oracle® Cloud Infrastructure (OCI) and on the Oracle® Cloud VMware Solution (OCVS) service, ensure volume performance units (VPUs) for your block volume are sized for your workloads.

The following list provides some guidance in selecting the VPUs needed for specific performance needs:

- Test or proof of concept environment: 100 GB, and 20 to 30 VPUs.
- Base-production environment: 500 GB, and 60 VPUs.
- Heavy-use production environment: More than 500 GB, and 100 or more VPUs.

Consider allocating additional VPUs to give enough capacity for updates and scaling activities. See [Block Volume Performance Levels \(Oracle documentation\)](#).

23.3.1.2. VMware vSphere CSI Driver Operator requirements

To install the vSphere Container Storage Interface (CSI) Driver Operator, the following requirements must be met:

- VMware vSphere version: 7.0 Update 2 or later; 8.0 Update 1 or later
- vCenter version: 7.0 Update 2 or later; 8.0 Update 1 or later
- Virtual machines of hardware version 15 or later
- No third-party vSphere CSI driver already installed in the cluster

If a third-party vSphere CSI driver is present in the cluster, OpenShift Container Platform does not overwrite it. The presence of a third-party vSphere CSI driver prevents OpenShift Container Platform from updating to OpenShift Container Platform 4.13 or later.



NOTE

The VMware vSphere CSI Driver Operator is supported only on clusters deployed with **platform: vsphere** in the installation manifest.

You can create a custom role for the Container Storage Interface (CSI) driver, the vSphere CSI Driver Operator, and the vSphere Problem Detector Operator. The custom role can include privilege sets that assign a minimum set of permissions to each vSphere object. This means that the CSI driver, the vSphere CSI Driver Operator, and the vSphere Problem Detector Operator can establish a basic interaction with these objects.



IMPORTANT

Installing an OpenShift Container Platform cluster in a vCenter is tested against a full list of privileges as described in the "Required vCenter account privileges" section. By adhering to the full list of privileges, you can reduce the possibility of unexpected and unsupported behaviors that might occur when creating a custom role with a set of restricted privileges.

Additional resources

- To remove a third-party vSphere CSI driver, see [Removing a third-party vSphere CSI Driver](#) .
- To update the hardware version for your vSphere nodes, see [Updating hardware on nodes running in vSphere](#).
- [Minimum permissions for the storage components](#)

23.3.1.3. Requirements for a cluster with user-provisioned infrastructure

For a cluster that contains user-provisioned infrastructure, you must deploy all of the required machines.

This section describes the requirements for deploying OpenShift Container Platform on user-provisioned infrastructure.

23.3.1.3.1. vCenter requirements

Before you install an OpenShift Container Platform cluster on your vCenter that uses infrastructure that you provided, you must prepare your environment.

Required vCenter account privileges

To install an OpenShift Container Platform cluster in a vCenter, your vSphere account must include privileges for reading and creating the required resources. Using an account that has global administrative privileges is the simplest way to access all of the necessary permissions.

Example 23.8. Roles and privileges required for installation in vSphere API

vSphere object for role	When required	Required privileges in vSphere API
vSphere vCenter	Always	Cns.Searchable InventoryService.Tagging.AttachTag InventoryService.Tagging.CreateCategory InventoryService.Tagging.CreateTag InventoryService.Tagging.DeleteCategory InventoryService.Tagging.DeleteTag InventoryService.Tagging.EditCategory InventoryService.Tagging.EditTag Sessions.ValidateSession StorageProfile.Update StorageProfile.View

vSphere object for role	When required	Required privileges in vSphere API
vSphere vCenter Cluster	If VMs will be created in the cluster root	Host.Config.Storage Resource.AssignVMToPool VApp.AssignResourcePool VApp.Import VirtualMachine.Config.AddNewDisk
vSphere vCenter Resource Pool	If an existing resource pool is provided	Resource.AssignVMToPool VApp.AssignResourcePool VApp.Import VirtualMachine.Config.AddNewDisk
vSphere Datastore	Always	Datastore.AllocateSpace Datastore.Browse Datastore.FileManagement InventoryService.Tagging.ObjectAttachable
vSphere Port Group	Always	Network.Assign
Virtual Machine Folder	Always	InventoryService.Tagging.ObjectAttachable Resource.AssignVMToPool VApp.Import VirtualMachine.Config.AddExistingDisk VirtualMachine.Config.AddNewDisk VirtualMachine.Config.AddRemoveDevice VirtualMachine.Config.AdvancedConfig VirtualMachine.Config.Annotation VirtualMachine.Config.CPUCount VirtualMachine.Config.DiskExtend VirtualMachine.Config.DiskLease VirtualMachine.Config.EditDevice VirtualMachine.Config.Memory VirtualMachine.Config.RemoveDisk VirtualMachine.Config.Rename Host.Config.Storage VirtualMachine.Config.Rese

vSphere object for role	When required	Required privileges in vSphere API
		VirtualMachine.Config.Settings VirtualMachine.Config.UpgradeVirtualHardware VirtualMachine.Interact.GuestControl VirtualMachine.Interact.PowerOff VirtualMachine.Interact.PowerOn VirtualMachine.Interact.Reset VirtualMachine.Inventory.Create VirtualMachine.Inventory.CreateFromExisting VirtualMachine.Inventory.Delete VirtualMachine.Provisioning.Clone VirtualMachine.Provisioning.MarkAsTemplate VirtualMachine.Provisioning.DeployTemplate
vSphere vCenter Datacenter	<p>If the installation program creates the virtual machine folder. For user-provisioned infrastructure,</p> <p>VirtualMachine.Inventory.Create and VirtualMachine.Inventory.Delete privileges are optional if your cluster does not use the Machine API. See the "Minimum permissions for the Machine API" table.</p>	InventoryService.Tagging.ObjectAttachable Resource.AssignVMToPool VirtualMachine.Config.AddExistingDisk VirtualMachine.Config.AddNewDisk VirtualMachine.Config.AddRemoveDevice VirtualMachine.Config.AdvancedConfig VirtualMachine.Config.Annotation VirtualMachine.Config.CPUCount VirtualMachine.Config.DiskExtend VirtualMachine.Config.DiskLease VirtualMachine.Config.EditDevice VirtualMachine.Config.Memory VirtualMachine.Config.RemoveDisk VirtualMachine.Config.Rename VirtualMachine.Config.Rese

vSphere object for role	When required	Required privileges in vSphere API
		VirtualMachine.Config.GuestInfo VirtualMachine.Config.Resource VirtualMachine.Config.Settings VirtualMachine.Config.UpgradeVirtualHardware VirtualMachine.Interact.GuestControl VirtualMachine.Interact.PowerOff VirtualMachine.Interact.PowerOn VirtualMachine.Interact.Reset VirtualMachine.Inventory.Create VirtualMachine.Inventory.CreateFromExisting VirtualMachine.Inventory.Delete VirtualMachine.Provisioning.Clone VirtualMachine.Provisioning.DeployTemplate VirtualMachine.Provisioning.MarkAsTemplate Folder.Create Folder.Delete

Example 23.9. Roles and privileges required for installation in vCenter graphical user interface (GUI)

vSphere object for role	When required	Required privileges in vCenter GUI

vSphere object for role	When required	Required privileges in vCenter GUI
vSphere vCenter	Always	Cns.Searchable "vSphere Tagging"."Assign or Unassign vSphere Tag" "vSphere Tagging"."Create vSphere Tag Category" "vSphere Tagging"."Create vSphere Tag" vSphere Tagging"."Delete vSphere Tag Category" "vSphere Tagging"."Delete vSphere Tag" "vSphere Tagging"."Edit vSphere Tag Category" "vSphere Tagging"."Edit vSphere Tag" Sessions."Validate session" "Profile-driven storage"."Profile-driven storage update" "Profile-driven storage"."Profile-driven storage view"
vSphere vCenter Cluster	If VMs will be created in the cluster root	Host.Configuration."Storage partition configuration" Resource."Assign virtual machine to resource pool" VApp."Assign resource pool" VApp.Import "Virtual machine"."Change Configuration"."Add new disk"
vSphere vCenter Resource Pool	If an existing resource pool is provided	Host.Configuration."Storage partition configuration" Resource."Assign virtual machine to resource pool" VApp."Assign resource pool" VApp.Import "Virtual machine"."Change Configuration"."Add new disk"
vSphere Datastore	Always	Datastore."Allocate space" Datastore."Browse datastore" Datastore."Low level file operations" "vSphere Tagging"."Assign or Unassign vSphere Tag on Object"

vSphere object for role	When required	Required privileges in vCenter GUI
vSphere Port Group	Always	Network."Assign network"
Virtual Machine Folder	Always	"vSphere Tagging"."Assign or Unassign vSphere Tag on Object" Resource."Assign virtual machine to resource pool" VApp.Import "Virtual machine"."Change Configuration"."Add existing disk" "Virtual machine"."Change Configuration"."Add new disk" "Virtual machine"."Change Configuration"."Add or remove device" "Virtual machine"."Change Configuration"."Advanced configuration" "Virtual machine"."Change Configuration"."Set annotation" "Virtual machine"."Change Configuration"."Change CPU count" "Virtual machine"."Change Configuration"."Extend virtual disk" "Virtual machine"."Change Configuration"."Acquire disk lease" "Virtual machine"."Change Configuration"."Modify device settings" "Virtual machine"."Change Configuration"."Change Memory" "Virtual machine"."Change Configuration"."Remove disk" "Virtual machine"."Change Configuration".Rename "Virtual machine"."Change Configuration"."Reset guest information" "Virtual machine"."Change Configuration"."Change resource" "Virtual machine"."Change Configuration"."Change Settings" "Virtual machine"."Change Configuration"."Upgrade virtual machine compatibility"

vSphere object for role	When required	Required privileges in vCenter GUI operating system management by VIX API"
		"Virtual machine".Interaction."Guest operating system management by VIX API" "Virtual machine".Interaction."Power off" "Virtual machine".Interaction."Power on" "Virtual machine".Interaction.Reset "Virtual machine"."Edit Inventory"."Create new" "Virtual machine"."Edit Inventory"."Create from existing" "Virtual machine"."Edit Inventory"."Remove" "Virtual machine".Provisioning."Clone virtual machine" "Virtual machine".Provisioning."Mark as template" "Virtual machine".Provisioning."Deploy template"
vSphere vCenter Datacenter	If the installation program creates the virtual machine folder. For user-provisioned infrastructure, VirtualMachine.Inventory.Create and VirtualMachine.Inventory.Delete privileges are optional if your cluster does not use the Machine API.	"vSphere Tagging"."Assign or Unassign vSphere Tag on Object" Resource."Assign virtual machine to resource pool" VApp.Import "Virtual machine"."Change Configuration"."Add existing disk" "Virtual machine"."Change Configuration"."Add new disk" "Virtual machine"."Change Configuration"."Add or remove device" "Virtual machine"."Change Configuration"."Advanced configuration" "Virtual machine"."Change Configuration"."Set annotation" "Virtual machine"."Change Configuration"."Change CPU count" "Virtual machine"."Change Configuration"."Extend virtual disk"

vSphere object for role	When required	"Virtual machine"."Change Configuration"."Acquire disk lease"
		"Virtual machine"."Change Configuration"."Modify device settings" "Virtual machine"."Change Configuration"."Change Memory" "Virtual machine"."Change Configuration"."Remove disk" "Virtual machine"."Change Configuration".Rename "Virtual machine"."Change Configuration"."Reset guest information" "Virtual machine"."Change Configuration"."Change resource" "Virtual machine"."Change Configuration"."Change Settings" "Virtual machine"."Change Configuration"."Upgrade virtual machine compatibility" "Virtual machine".Interaction."Guest operating system management by VIX API" "Virtual machine".Interaction."Power off" "Virtual machine".Interaction."Power on" "Virtual machine".Interaction.Reset "Virtual machine"."Edit Inventory"."Create new" "Virtual machine"."Edit Inventory"."Create from existing" "Virtual machine"."Edit Inventory"."Remove" "Virtual machine".Provisioning."Clone virtual machine" "Virtual machine".Provisioning."Deploy template" "Virtual machine".Provisioning."Mark as template" Folder."Create folder" Folder."Delete folder"

Additionally, the user requires some **ReadOnly** permissions, and some of the roles require permission to propagate the permissions to child objects. These settings vary depending on whether or not you install the cluster into an existing folder.

Example 23.10. Required permissions and propagation settings


vSphere object	When required	Propagate to children	Permissions required
vSphere vCenter	Always	False	Listed required privileges
vSphere vCenter Datacenter	Existing folder	False	ReadOnly permission
	Installation program creates the folder	True	Listed required privileges
vSphere vCenter Cluster	Existing resource pool	False	ReadOnly permission
	VMs in cluster root	True	Listed required privileges
vSphere vCenter Datastore	Always	False	Listed required privileges
vSphere Switch	Always	False	ReadOnly permission
vSphere Port Group	Always	False	Listed required privileges
vSphere vCenter Virtual Machine Folder	Existing folder	True	Listed required privileges
vSphere vCenter Resource Pool	Existing resource pool	True	Listed required privileges

For more information about creating an account with only the required privileges, see [vSphere Permissions and User Management Tasks](#) in the vSphere documentation.

Minimum required vCenter account privileges

After you create a custom role and assign privileges to it, you can create permissions by selecting specific vSphere objects and then assigning the custom role to a user or group for each object.

Before you create permissions or request for the creation of permissions for a vSphere object, determine what minimum permissions apply to the vSphere object. By doing this task, you can ensure a basic interaction exists between a vSphere object and OpenShift Container Platform architecture.



IMPORTANT

If you create a custom role and you do not assign privileges to it, the vSphere Server by default assigns a **Read Only** role to the custom role. Note that for the cloud provider API, the custom role only needs to inherit the privileges of the **Read Only** role.

Consider creating a custom role when an account with global administrative privileges does not meet your needs.



IMPORTANT

Accounts that are not configured with the required privileges are unsupported. Installing an OpenShift Container Platform cluster in a vCenter is tested against a full list of privileges as described in the "Required vCenter account privileges" section. By adhering to the full list of privileges, you can reduce the possibility of unexpected behaviors that might occur when creating a custom role with a restricted set of privileges.

The following tables list the minimum permissions for a vSphere object that interacts with specific OpenShift Container Platform architecture.

Example 23.11. Minimum permissions for post-installation management of components

vSphere object for role	When required	Required privileges
vSphere vCenter	Always	Cns.Searchable InventoryService.Tagging.AttachTag InventoryService.Tagging.CreateCategory InventoryService.Tagging.CreateTag InventoryService.Tagging.DeleteCategory InventoryService.Tagging.DeleteTag InventoryService.Tagging.EditCategory InventoryService.Tagging.EditTag Sessions.ValidateSession StorageProfile.Update StorageProfile.View
vSphere vCenter Cluster	If you intend to create VMs in the cluster root	Host.Config.StorageResource.AssignVMToolPool
vSphere vCenter Resource Pool	If you provide an existing resource pool in the install-config.yaml file	Host.Config.Storage

vSphere object for role	When required	Required privileges
vSphere Datastore	Always	Datastore.AllocateSpace Datastore.Browse Datastore.FileManagement InventoryService.Tagging.ObjectAttachable
vSphere Port Group	Always	Network.Assign
Virtual Machine Folder	Always	VirtualMachine.Config.AddExistingDisk VirtualMachine.Config.AddRemoveDevice VirtualMachine.Config.AdvancedConfig VirtualMachine.Config.Annotation VirtualMachine.Config.CPUCount VirtualMachine.Config.DiskExtend VirtualMachine.Config.Memory VirtualMachine.Config.Settings VirtualMachine.Interact.PowerOff VirtualMachine.Interact.PowerOn VirtualMachine.Inventory.CreateFromExisting VirtualMachine.Inventory.Delete VirtualMachine.Provisioning.Clone VirtualMachine.Provisioning.DeployTemplate

vSphere object for role	When required	Required privileges
vSphere vCenter Datacenter	If the installation program creates the virtual machine folder. For user-provisioned infrastructure, VirtualMachine.Inventory.Create and VirtualMachine.Inventory.Delete privileges are optional if your cluster does not use the Machine API. If your cluster does use the Machine API and you want to set the minimum set of permissions for the API, see the "Minimum permissions for the Machine API" table.	Resource.AssignVMT oPool VirtualMachine.Config .AddExistingDisk VirtualMachine.Config .AddRemoveDevice VirtualMachine.Interac t.PowerOff VirtualMachine.Interac t.PowerOn VirtualMachine.Provisi oning.DeployTemplat e

Example 23.12. Minimum permissions for the storage components

vSphere object for role	When required	Required privileges
vSphere vCenter	Always	Cns.Searchable InventoryService.Tag ging.CreateCategory InventoryService.Tag ging.CreateTag InventoryService.Tag ging.EditCategory InventoryService.Tag ging.EditTag StorageProfile.Update StorageProfile.View
vSphere vCenter Cluster	If you intend to create VMs in the cluster root	Host.Config.Storage
vSphere vCenter Resource Pool	If you provide an existing resource pool in the install-config.yaml file	Host.Config.Storage
vSphere Datastore	Always	Datastore.Browse Datastore.FileManage ment InventoryService.Tag ging.ObjectAttachable
vSphere Port Group	Always	Read Only

vSphere object for role	When required	Required privileges
Virtual Machine Folder	Always	VirtualMachine.Config.AddExistingDisk VirtualMachine.Config.AddRemoveDevice
vSphere vCenter Datacenter	If the installation program creates the virtual machine folder. For user-provisioned infrastructure, VirtualMachine.Inventory.Create and VirtualMachine.Inventory.Delete privileges are optional if your cluster does not use the Machine API. If your cluster does use the Machine API and you want to set the minimum set of permissions for the API, see the "Minimum permissions for the Machine API" table.	VirtualMachine.Config.AddExistingDisk VirtualMachine.Config.AddRemoveDevice

Example 23.13. Minimum permissions for the Machine API

vSphere object for role	When required	Required privileges
vSphere vCenter	Always	InventoryService.Tagging.AttachTag InventoryService.Tagging.CreateCategory InventoryService.Tagging.CreateTag InventoryService.Tagging.DeleteCategory InventoryService.Tagging.DeleteTag InventoryService.Tagging.EditCategory InventoryService.Tagging.EditTag Sessions.ValidateSession StorageProfile.Update StorageProfile.View
vSphere vCenter Cluster	If you intend to create VMs in the cluster root	Resource.AssignVMT oPool
vSphere vCenter Resource Pool	If you provide an existing resource pool in the install-config.yaml file	Read Only

vSphere object for role	When required	Required privileges
vSphere Datastore	Always	Datastore.AllocateSpace Datastore.Browse
vSphere Port Group	Always	Network.Assign
Virtual Machine Folder	Always	VirtualMachine.Config.AddRemoveDevice VirtualMachine.Config.AdvancedConfig VirtualMachine.Config.Annotation VirtualMachine.Config.CPUCount VirtualMachine.Config.DiskExtend VirtualMachine.Config.Memory VirtualMachine.Config.Settings VirtualMachine.Interact.PowerOff VirtualMachine.Interact.PowerOn VirtualMachine.Inventory.CreateFromExisting VirtualMachine.Inventory.Delete VirtualMachine.Provisioning.Clone VirtualMachine.Provisioning.DeployTemplate
vSphere vCenter Datacenter	If the installation program creates the virtual machine folder. For user-provisioned infrastructure, VirtualMachine.Inventory.Create and VirtualMachine.Inventory.Delete privileges are optional if your cluster does not use the Machine API.	Resource.AssignVMT oPool VirtualMachine.Interact.PowerOff VirtualMachine.Interact.PowerOn VirtualMachine.Provisioning.DeployTemplate


Using OpenShift Container Platform with vMotion

If you intend on using vMotion in your vSphere environment, consider the following before installing an OpenShift Container Platform cluster.

- OpenShift Container Platform generally supports compute-only vMotion, where *generally* implies that you meet all VMware best practices for vMotion. To help ensure the uptime of your compute and control plane nodes, ensure that you follow the VMware best practices for vMotion, and use VMware anti-affinity rules to improve the availability of OpenShift Container Platform during maintenance or hardware issues.

For more information about vMotion and anti-affinity rules, see the VMware vSphere documentation for [vMotion networking requirements](#) and [VM anti-affinity rules](#).

- Using Storage vMotion can cause issues and is not supported. If you are using vSphere volumes in your pods, migrating a VM across datastores, either manually or through Storage vMotion, causes invalid references within OpenShift Container Platform persistent volume (PV) objects that can result in data loss.
- OpenShift Container Platform does not support selective migration of VMDKs across datastores, using datastore clusters for VM provisioning or for dynamic or static provisioning of PVs, or using a datastore that is part of a datastore cluster for dynamic or static provisioning of PVs.



IMPORTANT

You can specify the path of any datastore that exists in a datastore cluster. By default, Storage Distributed Resource Scheduler (SDRS), which uses Storage vMotion, is automatically enabled for a datastore cluster. Red Hat does not support Storage vMotion, so you must disable Storage DRS to avoid data loss issues for your OpenShift Container Platform cluster.

If you must specify VMs across multiple datastores, use a **datastore** object to specify a failure domain in your cluster's **install-config.yaml** configuration file. For more information, see "VMware vSphere region and zone enablement".

Cluster resources

When you deploy an OpenShift Container Platform cluster that uses infrastructure that you provided, you must create the following resources in your vCenter instance:

- 1 Folder
- 1 Tag category
- 1 Tag
- Virtual machines:
 - 1 template
 - 1 temporary bootstrap node
 - 3 control plane nodes
 - 3 compute machines

Although these resources use 856 GB of storage, the bootstrap node is destroyed during the cluster installation process. A minimum of 800 GB of storage is required to use a standard cluster.

If you deploy more compute machines, the OpenShift Container Platform cluster will use more storage.

Cluster limits

Available resources vary between clusters. The number of possible clusters within a vCenter is limited primarily by available storage space and any limitations on the number of required resources. Be sure to consider both limitations to the vCenter resources that the cluster creates and the resources that you require to deploy a cluster, such as IP addresses and networks.

Networking requirements

Use Dynamic Host Configuration Protocol (DHCP) for the network and ensure that the DHCP server is configured to provide persistent IP addresses to the cluster machines.



NOTE


You do not need to use the DHCP for the network if you want to provision nodes with static IP addresses.

Configure the default gateway to use the DHCP server. All nodes must be in the same VLAN. You cannot scale the cluster using a second VLAN as a Day 2 operation.

You must use the Dynamic Host Configuration Protocol (DHCP) for the network and ensure that the DHCP server is configured to provide persistent IP addresses to the cluster machines. In the DHCP lease, you must configure the DHCP to use the default gateway. All nodes must be in the same VLAN. You cannot scale the cluster using a second VLAN as a Day 2 operation.

If you are installing to a restricted environment, the VM in your restricted network must have access to vCenter so that it can provision and manage nodes, persistent volume claims (PVCs), and other resources.

Additionally, you must create the following networking resources before you install the OpenShift Container Platform cluster:



NOTE

It is recommended that each OpenShift Container Platform node in the cluster must have access to a Network Time Protocol (NTP) server that is discoverable via DHCP. Installation is possible without an NTP server. However, asynchronous server clocks will cause errors, which NTP server prevents.

DNS records

You must create DNS records for two static IP addresses in the appropriate DNS server for the vCenter instance that hosts your OpenShift Container Platform cluster. In each record, **<cluster_name>** is the cluster name and **<base_domain>** is the cluster base domain that you specify when you install the cluster. A complete DNS record takes the form: **<component>.<cluster_name>.<base_domain>..**

Table 23.20. Required DNS records

Component	Record	Description
-----------	--------	-------------

Component	Record	Description
API VIP	api.<cluster_name>.<base_domain>.	This DNS A/AAAA or CNAME (Canonical Name) record must point to the load balancer for the control plane machines. This record must be resolvable by both clients external to the cluster and from all the nodes within the cluster.
Ingress VIP	*.apps.<cluster_name>.<base_domain>.	A wildcard DNS A/AAAA or CNAME record that points to the load balancer that targets the machines that run the Ingress router pods, which are the worker nodes by default. This record must be resolvable by both clients external to the cluster and from all the nodes within the cluster.

Additional resources

- [Creating a compute machine set on vSphere](#)

23.3.1.3.2. Required machines for cluster installation

The smallest OpenShift Container Platform clusters require the following hosts:

Table 23.21. Minimum required hosts

Hosts	Description
One temporary bootstrap machine	The cluster requires the bootstrap machine to deploy the OpenShift Container Platform cluster on the three control plane machines. You can remove the bootstrap machine after you install the cluster.
Three control plane machines	The control plane machines run the Kubernetes and OpenShift Container Platform services that form the control plane.
At least two compute machines, which are also known as worker machines.	The workloads requested by OpenShift Container Platform users run on the compute machines.



IMPORTANT

To maintain high availability of your cluster, use separate physical hosts for these cluster machines.

The bootstrap and control plane machines must use Red Hat Enterprise Linux CoreOS (RHCOS) as the operating system. However, the compute machines can choose between Red Hat Enterprise Linux CoreOS (RHCOS), Red Hat Enterprise Linux (RHEL) 8.6 and later.

Note that RHCOS is based on Red Hat Enterprise Linux (RHEL) 9.2 and inherits all of its hardware certifications and requirements. See [Red Hat Enterprise Linux technology capabilities and limits](#) .

23.3.1.3.3. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 23.22. Minimum resource requirements

Machine	Operating System	vCPU	Virtual RAM	Storage	Input/Output Per Second (IOPS)[1]
Bootstrap	RHCOS	4	16 GB	100 GB	300
Control plane	RHCOS	4	16 GB	100 GB	300
Compute	RHCOS, RHEL 8.6 and later [2]	2	8 GB	100 GB	300

1. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes which require a 10 ms p99 fsync duration. Note that on many cloud platforms, storage size and IOPS scale together, so you might need to over-allocate storage volume to obtain sufficient performance.
2. As with all user-provisioned installations, if you choose to use RHEL compute machines in your cluster, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Use of RHEL 7 compute machines is deprecated and has been removed in OpenShift Container Platform 4.10 and later.

**NOTE**

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

Additional resources

- [Optimizing storage](#)

23.3.1.3.4. Requirements for encrypting virtual machines

You can encrypt your virtual machines prior to installing OpenShift Container Platform 4.16 by meeting the following requirements.

- You have configured a Standard key provider in vSphere. For more information, see [Adding a KMS to vCenter Server](#).

**IMPORTANT**

The Native key provider in vCenter is not supported. For more information, see [vSphere Native Key Provider Overview](#).

- You have enabled host encryption mode on all of the ESXi hosts that are hosting the cluster. For more information, see [Enabling host encryption mode](#).
- You have a vSphere account which has all cryptographic privileges enabled. For more information, see [Cryptographic Operations Privileges](#).

When you deploy the OVF template in the section titled "Installing RHCOS and starting the OpenShift Container Platform bootstrap process", select the option to "Encrypt this virtual machine" when you are selecting storage for the OVF template. After completing cluster installation, create a storage class that uses the encryption storage policy you used to encrypt the virtual machines.

Additional resources

- [Creating an encrypted storage class](#)

23.3.1.3.5. Certificate signing requests management

Because your cluster has limited access to automatic machine management when you use infrastructure that you provision, you must provide a mechanism for approving cluster certificate signing requests

(CSRs) after installation. The **kube-controller-manager** only approves the kubelet client CSRs. The **machine-approver** cannot guarantee the validity of a serving certificate that is requested by using kubelet credentials because it cannot confirm that the correct machine issued the request. You must determine and implement a method of verifying the validity of the kubelet serving certificate requests and approving them.

23.3.1.3.6. Networking requirements for user-provisioned infrastructure

All the Red Hat Enterprise Linux CoreOS (RHCOS) machines require networking to be configured in **initramfs** during boot to fetch their Ignition config files.

During the initial boot, the machines require an IP address configuration that is set either through a DHCP server or statically by providing the required boot options. After a network connection is established, the machines download their Ignition config files from an HTTP or HTTPS server. The Ignition config files are then used to set the exact state of each machine. The Machine Config Operator completes more changes to the machines, such as the application of new certificates or keys, after installation.

It is recommended to use a DHCP server for long-term management of the cluster machines. Ensure that the DHCP server is configured to provide persistent IP addresses, DNS server information, and hostnames to the cluster machines.



NOTE

If a DHCP service is not available for your user-provisioned infrastructure, you can instead provide the IP networking configuration and the address of the DNS server to the nodes at RHCOS install time. These can be passed as boot arguments if you are installing from an ISO image. See the *Installing RHCOS and starting the OpenShift Container Platform bootstrap process* section for more information about static IP provisioning and advanced networking options.

The Kubernetes API server must be able to resolve the node names of the cluster machines. If the API servers and worker nodes are in different zones, you can configure a default DNS search zone to allow the API server to resolve the node names. Another supported approach is to always refer to hosts by their fully-qualified domain names in both the node objects and all DNS requests.

23.3.1.3.6.1. Setting the cluster node hostnames through DHCP

On Red Hat Enterprise Linux CoreOS (RHCOS) machines, the hostname is set through NetworkManager. By default, the machines obtain their hostname through DHCP. If the hostname is not provided by DHCP, set statically through kernel arguments, or another method, it is obtained through a reverse DNS lookup. Reverse DNS lookup occurs after the network has been initialized on a node and can take time to resolve. Other system services can start prior to this and detect the hostname as **localhost** or similar. You can avoid this by using DHCP to provide the hostname for each cluster node.

Additionally, setting the hostnames through DHCP can bypass any manual DNS record name configuration errors in environments that have a DNS split-horizon implementation.

23.3.1.3.6.2. Network connectivity requirements

You must configure the network connectivity between machines to allow OpenShift Container Platform cluster components to communicate. Each machine must be able to resolve the hostnames of all other machines in the cluster.

This section provides details about the ports that are required.



IMPORTANT

In connected OpenShift Container Platform environments, all nodes are required to have internet access to pull images for platform containers and provide telemetry data to Red Hat.

Table 23.23. Ports used for all-machine to all-machine communications

Protocol	Port	Description
ICMP	N/A	Network reachability tests
TCP	1936	Metrics
	9000-9999	Host level services, including the node exporter on ports 9100-9101 and the Cluster Version Operator on port 9099 .
	10250-10259	The default ports that Kubernetes reserves
UDP	4789	VXLAN
	6081	Geneve
	9000-9999	Host level services, including the node exporter on ports 9100-9101 .
	500	IPsec IKE packets
	4500	IPsec NAT-T packets
	123	Network Time Protocol (NTP) on UDP port 123 If an external NTP time server is configured, you must open UDP port 123 .
TCP/UDP	30000-32767	Kubernetes node port
ESP	N/A	IPsec Encapsulating Security Payload (ESP)

Table 23.24. Ports used for all-machine to control plane communications

Protocol	Port	Description
TCP	6443	Kubernetes API

Table 23.25. Ports used for control plane machine to control plane machine communications

Protocol	Port	Description
TCP	2379-2380	etcd server and peer ports

Ethernet adaptor hardware address requirements

When provisioning VMs for the cluster, the ethernet interfaces configured for each VM must use a MAC address from the VMware Organizationally Unique Identifier (OUI) allocation ranges:

- **00:05:69:00:00:00 to 00:05:69:FF:FF:FF**
- **00:0c:29:00:00:00 to 00:0c:29:FF:FF:FF**
- **00:1c:14:00:00:00 to 00:1c:14:FF:FF:FF**
- **00:50:56:00:00:00 to 00:50:56:3F:FF:FF**

If a MAC address outside the VMware OUI is used, the cluster installation will not succeed.

NTP configuration for user-provisioned infrastructure

OpenShift Container Platform clusters are configured to use a public Network Time Protocol (NTP) server by default. If you want to use a local enterprise NTP server, or if your cluster is being deployed in a disconnected network, you can configure the cluster to use a specific time server. For more information, see the documentation for *Configuring chrony time service*.

If a DHCP server provides NTP server information, the chrony time service on the Red Hat Enterprise Linux CoreOS (RHCOS) machines read the information and can sync the clock with the NTP servers.

Additional resources

- [Configuring chrony time service](#)

23.3.1.3.7. User-provisioned DNS requirements

In OpenShift Container Platform deployments, DNS name resolution is required for the following components:

- The Kubernetes API
- The OpenShift Container Platform application wildcard
- The bootstrap, control plane, and compute machines

Reverse DNS resolution is also required for the Kubernetes API, the bootstrap machine, the control plane machines, and the compute machines.


DNS A/AAAA or CNAME records are used for name resolution and PTR records are used for reverse name resolution. The reverse records are important because Red Hat Enterprise Linux CoreOS (RHCOS) uses the reverse records to set the hostnames for all the nodes, unless the hostnames are provided by DHCP. Additionally, the reverse records are used to generate the certificate signing requests (CSR) that OpenShift Container Platform needs to operate.

**NOTE**

It is recommended to use a DHCP server to provide the hostnames to each cluster node. See the *DHCP recommendations for user-provisioned infrastructure* section for more information.

The following DNS records are required for a user-provisioned OpenShift Container Platform cluster and they must be in place before installation. In each record, **<cluster_name>** is the cluster name and **<base_domain>** is the base domain that you specify in the **install-config.yaml** file. A complete DNS record takes the form: **<component>.<cluster_name>.<base_domain>.**

Table 23.26. Required DNS records

Component	Record	Description
Kubernetes API	api.<cluster_name>.<base_domain>.	A DNS A/AAAA or CNAME record, and a DNS PTR record, to identify the API load balancer. These records must be resolvable by both clients external to the cluster and from all the nodes within the cluster.
	api-int.<cluster_name>.<base_domain>.	A DNS A/AAAA or CNAME record, and a DNS PTR record, to internally identify the API load balancer. These records must be resolvable from all the nodes within the cluster.
		 <p>IMPORTANT</p> <p>The API server must be able to resolve the worker nodes by the hostnames that are recorded in Kubernetes. If the API server cannot resolve the node names, then proxied API calls can fail, and you cannot retrieve logs from pods.</p>
Routes	*.apps.<cluster_name>.<base_domain>.	<p>A wildcard DNS A/AAAA or CNAME record that refers to the application ingress load balancer. The application ingress load balancer targets the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default. These records must be resolvable by both clients external to the cluster and from all the nodes within the cluster.</p> <p>For example, console-openshift-console.apps.<cluster_name>.<base_domain> is used as a wildcard route to the OpenShift Container Platform console.</p>
Bootstrap machine	bootstrap.<cluster_name>.<base_domain>.	A DNS A/AAAA or CNAME record, and a DNS PTR record, to identify the bootstrap machine. These records must be resolvable by the nodes within the cluster.

Component	Record	Description
Control plane machines	<control_plane><n>. <cluster_name>. <base_domain>.	DNS A/AAAA or CNAME records and DNS PTR records to identify each machine for the control plane nodes. These records must be resolvable by the nodes within the cluster.
Compute machines	<compute><n>. <cluster_name>. <base_domain>.	DNS A/AAAA or CNAME records and DNS PTR records to identify each machine for the worker nodes. These records must be resolvable by the nodes within the cluster.



NOTE

In OpenShift Container Platform 4.4 and later, you do not need to specify etcd host and SRV records in your DNS configuration.

TIP

You can use the **dig** command to verify name and reverse name resolution. See the section on *Validating DNS resolution for user-provisioned infrastructure* for detailed validation steps.

23.3.1.3.7.1. Example DNS configuration for user-provisioned clusters

This section provides A and PTR record configuration samples that meet the DNS requirements for deploying OpenShift Container Platform on user-provisioned infrastructure. The samples are not meant to provide advice for choosing one DNS solution over another.

In the examples, the cluster name is **ocp4** and the base domain is **example.com**.

Example DNS A record configuration for a user-provisioned cluster

The following example is a BIND zone file that shows sample A records for name resolution in a user-provisioned cluster.

Example 23.14. Sample DNS zone database

```
$TTL 1W
@ IN SOA ns1.example.com. root (
    2019070700 ; serial
    3H ; refresh (3 hours)
    30M ; retry (30 minutes)
    2W ; expiry (2 weeks)
    1W ) ; minimum (1 week)
IN NS ns1.example.com.
IN MX 10 smtp.example.com.
;
;
ns1.example.com. IN A 192.168.1.5
smtp.example.com. IN A 192.168.1.5
;
helper.example.com. IN A 192.168.1.5
```

```

helper.ocp4.example.com. IN A 192.168.1.5
;
api.ocp4.example.com. IN A 192.168.1.5 1
api-int.ocp4.example.com. IN A 192.168.1.5 2
;
*.apps.ocp4.example.com. IN A 192.168.1.5 3
;
bootstrap.ocp4.example.com. IN A 192.168.1.96 4
;
control-plane0.ocp4.example.com. IN A 192.168.1.97 5
control-plane1.ocp4.example.com. IN A 192.168.1.98 6
control-plane2.ocp4.example.com. IN A 192.168.1.99 7
;
compute0.ocp4.example.com. IN A 192.168.1.11 8
compute1.ocp4.example.com. IN A 192.168.1.7 9
;
;EOF

```

- 1 Provides name resolution for the Kubernetes API. The record refers to the IP address of the API load balancer.
- 2 Provides name resolution for the Kubernetes API. The record refers to the IP address of the API load balancer and is used for internal cluster communications.
- 3 Provides name resolution for the wildcard routes. The record refers to the IP address of the application ingress load balancer. The application ingress load balancer targets the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.



NOTE

In the example, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

- 4 Provides name resolution for the bootstrap machine.
- 5 6 7 Provides name resolution for the control plane machines.
- 8 9 Provides name resolution for the compute machines.

Example DNS PTR record configuration for a user-provisioned cluster

The following example BIND zone file shows sample PTR records for reverse name resolution in a user-provisioned cluster.

Example 23.15. Sample DNS zone database for reverse records

```

$TTL 1W
@ IN SOA ns1.example.com. root (
  2019070700 ; serial

```

```

3H ; refresh (3 hours)
30M ; retry (30 minutes)
2W ; expiry (2 weeks)
1W ) ; minimum (1 week)
IN NS ns1.example.com.
;
5.1.168.192.in-addr.arpa. IN PTR api.ocp4.example.com. 1
5.1.168.192.in-addr.arpa. IN PTR api-int.ocp4.example.com. 2
;
96.1.168.192.in-addr.arpa. IN PTR bootstrap.ocp4.example.com. 3
;
97.1.168.192.in-addr.arpa. IN PTR control-plane0.ocp4.example.com. 4
98.1.168.192.in-addr.arpa. IN PTR control-plane1.ocp4.example.com. 5
99.1.168.192.in-addr.arpa. IN PTR control-plane2.ocp4.example.com. 6
;
11.1.168.192.in-addr.arpa. IN PTR compute0.ocp4.example.com. 7
7.1.168.192.in-addr.arpa. IN PTR compute1.ocp4.example.com. 8
;
;EOF

```

- 1 Provides reverse DNS resolution for the Kubernetes API. The PTR record refers to the record name of the API load balancer.
- 2 Provides reverse DNS resolution for the Kubernetes API. The PTR record refers to the record name of the API load balancer and is used for internal cluster communications.
- 3 Provides reverse DNS resolution for the bootstrap machine.
- 4 5 6 Provides reverse DNS resolution for the control plane machines.
- 7 8 Provides reverse DNS resolution for the compute machines.




NOTE

A PTR record is not required for the OpenShift Container Platform application wildcard.

23.3.1.3.8. Load balancing requirements for user-provisioned infrastructure

Before you install OpenShift Container Platform, you must provision the API and application Ingress load balancing infrastructure. In production scenarios, you can deploy the API and application Ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.




NOTE

If you want to deploy the API and application Ingress load balancers with a Red Hat Enterprise Linux (RHEL) instance, you must purchase the RHEL subscription separately.

The load balancing infrastructure must meet the following requirements:

1. **API load balancer.** Provides a common endpoint for users, both human and machine, to interact with and configure the platform. Configure the following conditions:
 - Layer 4 load balancing only. This can be referred to as Raw TCP or SSL Passthrough mode.

- Layer 4 load balancing only. This can be referred to as Raw TCP or SSL Passthrough mode.
- A stateless load balancing algorithm. The options vary based on the load balancer implementation.



IMPORTANT

Do not configure session persistence for an API load balancer. Configuring session persistence for a Kubernetes API server might cause performance issues from excess application traffic for your OpenShift Container Platform cluster and the Kubernetes API that runs inside the cluster.

Configure the following ports on both the front and back of the load balancers:

Table 23.27. API load balancer

Port	Back-end machines (pool members)	Internal	External	Description
6443	Bootstrap and control plane. You remove the bootstrap machine from the load balancer after the bootstrap machine initializes the cluster control plane. You must configure the /readyz endpoint for the API server health check probe.	X	X	Kubernetes API server
22623	Bootstrap and control plane. You remove the bootstrap machine from the load balancer after the bootstrap machine initializes the cluster control plane.	X		Machine config server



NOTE

The load balancer must be configured to take a maximum of 30 seconds from the time the API server turns off the **/readyz** endpoint to the removal of the API server instance from the pool. Within the time frame after **/readyz** returns an error or becomes healthy, the endpoint must have been removed or added. Probing every 5 or 10 seconds, with two successful requests to become healthy and three to become unhealthy, are well-tested values.

2. **Application Ingress load balancer:** Provides an ingress point for application traffic flowing in from outside the cluster. A working configuration for the Ingress router is required for an OpenShift Container Platform cluster.

Configure the following conditions:

- Layer 4 load balancing only. This can be referred to as Raw TCP or SSL Passthrough mode.
- A connection-based or session-based persistence is recommended, based on the options available and types of applications that will be hosted on the platform.

TIP

If the true IP address of the client can be seen by the application Ingress load balancer, enabling source IP-based session persistence can improve performance for applications that use end-to-end TLS encryption.

Configure the following ports on both the front and back of the load balancers:

Table 23.28. Application Ingress load balancer

Port	Back-end machines (pool members)	Internal	External	Description
443	The machines that run the Ingress Controller pods, compute, or worker, by default.	X	X	HTTPS traffic
80	The machines that run the Ingress Controller pods, compute, or worker, by default.	X	X	HTTP traffic

**NOTE**

If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application Ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes.

23.3.1.3.8.1. Example load balancer configuration for user-provisioned clusters

This section provides an example API and application Ingress load balancer configuration that meets the load balancing requirements for user-provisioned clusters. The sample is an `/etc/haproxy/haproxy.cfg` configuration for an HAProxy load balancer. The example is not meant to provide advice for choosing one load balancing solution over another.

In the example, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

**NOTE**

If you are using HAProxy as a load balancer and SELinux is set to **enforcing**, you must ensure that the HAProxy service can bind to the configured TCP port by running **setsebool -P haproxy_connect_any=1**.

Example 23.16. Sample API and application Ingress load balancer configuration

```
global
log      127.0.0.1 local2
pidfile  /var/run/haproxy.pid
maxconn 4000
daemon
defaults
```


```

mode          http
log           global
option       dontlognull
option http-server-close
option       redispatch
retries      3
timeout http-request 10s
timeout queue 1m
timeout connect 10s
timeout client 1m
timeout server 1m
timeout http-keep-alive 10s
timeout check 10s
maxconn      3000
listen api-server-6443 1
bind *:6443
mode tcp
option httpchk GET /readyz HTTP/1.0
option log-health-checks
balance roundrobin
server bootstrap bootstrap.ocp4.example.com:6443 verify none check check-ssl inter 10s fall 2
rise 3 backup 2
server master0 master0.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
fall 2 rise 3
server master1 master1.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
fall 2 rise 3
server master2 master2.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
fall 2 rise 3
listen machine-config-server-22623 3
bind *:22623
mode tcp
server bootstrap bootstrap.ocp4.example.com:22623 check inter 1s backup 4
server master0 master0.ocp4.example.com:22623 check inter 1s
server master1 master1.ocp4.example.com:22623 check inter 1s
server master2 master2.ocp4.example.com:22623 check inter 1s
listen ingress-router-443 5
bind *:443
mode tcp
balance source
server compute0 compute0.ocp4.example.com:443 check inter 1s
server compute1 compute1.ocp4.example.com:443 check inter 1s
listen ingress-router-80 6
bind *:80
mode tcp
balance source
server compute0 compute0.ocp4.example.com:80 check inter 1s
server compute1 compute1.ocp4.example.com:80 check inter 1s

```

- 1** Port **6443** handles the Kubernetes API traffic and points to the control plane machines.
- 2** **4** The bootstrap entries must be in place before the OpenShift Container Platform cluster installation and they must be removed after the bootstrap process is complete.
- 3** Port **22623** handles the machine config server traffic and points to the control plane machines.

- 5 Port **443** handles the HTTPS traffic and points to the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.
- 6 Port **80** handles the HTTP traffic and points to the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.

**NOTE**

If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application Ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes.


TIP

If you are using HAProxy as a load balancer, you can check that the **haproxy** process is listening on ports **6443**, **22623**, **443**, and **80** by running **netstat -nltpu** on the HAProxy node.

23.3.2. Preparing to install a cluster using user-provisioned infrastructure

You prepare to install an OpenShift Container Platform cluster on vSphere by completing the following steps:

- Downloading the installation program.

**NOTE**

If you are installing in a disconnected environment, you extract the installation program from the mirrored content. For more information, see [Mirroring images for a disconnected installation](#).

- Installing the OpenShift CLI (**oc**).

**NOTE**

If you are installing in a disconnected environment, install **oc** to the mirror host.

- Generating an SSH key pair. You can use this key pair to authenticate into the OpenShift Container Platform cluster's nodes after it is deployed.
- Preparing the user-provisioned infrastructure.
- Validating DNS resolution.

23.3.2.1. Obtaining the installation program

Before you install OpenShift Container Platform, download the installation file on the host you are using for installation.

Prerequisites

- You have a computer that runs Linux or macOS, with 500 MB of local disk space.


Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select your infrastructure provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.



IMPORTANT

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.



IMPORTANT

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.

4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#) . This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

23.3.2.2. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.



IMPORTANT

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**. To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**. To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.

**NOTE**

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your PATH.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:


```
$ oc <command>
```

23.3.2.3. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.

**IMPORTANT**

Do not skip this procedure in production environments, where disaster recovery and debugging is required.

**NOTE**

You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.



NOTE

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.



NOTE

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```



NOTE

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program. If you install a cluster on infrastructure that you provision, you must provide the key to the installation program.

23.3.2.4. Preparing the user-provisioned infrastructure

Before you install OpenShift Container Platform on user-provisioned infrastructure, you must prepare the underlying infrastructure.

This section provides details about the high-level steps required to set up your cluster infrastructure in preparation for an OpenShift Container Platform installation. This includes configuring IP networking and network connectivity for your cluster nodes, enabling the required ports through your firewall, and setting up the required DNS and load balancing infrastructure.

After preparation, your cluster infrastructure must meet the requirements outlined in the *Requirements for a cluster with user-provisioned infrastructure* section.

Prerequisites

- You have reviewed the [OpenShift Container Platform 4.x Tested Integrations](#) page.
- You have reviewed the infrastructure requirements detailed in the *Requirements for a cluster with user-provisioned infrastructure* section.

Procedure

- If you are using DHCP to provide the IP networking configuration to your cluster nodes, configure your DHCP service.
 - Add persistent IP addresses for the nodes to your DHCP server configuration. In your configuration, match the MAC address of the relevant network interface to the intended IP address for each node.
 - When you use DHCP to configure IP addressing for the cluster machines, the machines also obtain the DNS server information through DHCP. Define the persistent DNS server address that is used by the cluster nodes through your DHCP server configuration.

**NOTE**

If you are not using a DHCP service, you must provide the IP networking configuration and the address of the DNS server to the nodes at RHCOS install time. These can be passed as boot arguments if you are installing from an ISO image. See the *Installing RHCOS and starting the OpenShift Container Platform bootstrap process* section for more information about static IP provisioning and advanced networking options.

- c. Define the hostnames of your cluster nodes in your DHCP server configuration. See the *Setting the cluster node hostnames through DHCP* section for details about hostname considerations.

**NOTE**

If you are not using a DHCP service, the cluster nodes obtain their hostname through a reverse DNS lookup.

2. Ensure that your network infrastructure provides the required network connectivity between the cluster components. See the *Networking requirements for user-provisioned infrastructure* section for details about the requirements.
3. Configure your firewall to enable the ports required for the OpenShift Container Platform cluster components to communicate. See *Networking requirements for user-provisioned infrastructure* section for details about the ports that are required.

**IMPORTANT**

By default, port **1936** is accessible for an OpenShift Container Platform cluster, because each control plane node needs access to this port.

Avoid using the Ingress load balancer to expose this port, because doing so might result in the exposure of sensitive information, such as statistics and metrics, related to Ingress Controllers.

4. Setup the required DNS infrastructure for your cluster.
 - a. Configure DNS name resolution for the Kubernetes API, the application wildcard, the bootstrap machine, the control plane machines, and the compute machines.
 - b. Configure reverse DNS resolution for the Kubernetes API, the bootstrap machine, the control plane machines, and the compute machines.
See the *User-provisioned DNS requirements* section for more information about the OpenShift Container Platform DNS requirements.
5. Validate your DNS configuration.
 - a. From your installation node, run DNS lookups against the record names of the Kubernetes API, the wildcard routes, and the cluster nodes. Validate that the IP addresses in the responses correspond to the correct components.
 - b. From your installation node, run reverse DNS lookups against the IP addresses of the load balancer and the cluster nodes. Validate that the record names in the responses correspond to the correct components.

See the *Validating DNS resolution for user-provisioned infrastructure* section for detailed DNS validation steps.

6. Provision the required API and application ingress load balancing infrastructure. See the *Load balancing requirements for user-provisioned infrastructure* section for more information about the requirements.




NOTE

Some load balancing solutions require the DNS name resolution for the cluster nodes to be in place before the load balancing is initialized.

23.3.2.5. Validating DNS resolution for user-provisioned infrastructure

You can validate your DNS configuration before installing OpenShift Container Platform on user-provisioned infrastructure.



IMPORTANT

The validation steps detailed in this section must succeed before you install your cluster.

Prerequisites

- You have configured the required DNS records for your user-provisioned infrastructure.

Procedure

1. From your installation node, run DNS lookups against the record names of the Kubernetes API, the wildcard routes, and the cluster nodes. Validate that the IP addresses contained in the responses correspond to the correct components.
 - a. Perform a lookup against the Kubernetes API record name. Check that the result points to the IP address of the API load balancer:

```
$ dig +noall +answer @<nameserver_ip> api.<cluster_name>.<base_domain> 1
```

- 1** Replace **<nameserver_ip>** with the IP address of the nameserver, **<cluster_name>** with your cluster name, and **<base_domain>** with your base domain name.

Example output

```
api.ocp4.example.com. 604800 IN A 192.168.1.5
```

- b. Perform a lookup against the Kubernetes internal API record name. Check that the result points to the IP address of the API load balancer:

```
$ dig +noall +answer @<nameserver_ip> api-int.<cluster_name>.<base_domain>
```

Example output


```
api-int.ocp4.example.com. 604800 IN A 192.168.1.5
```

- c. Test an example `*.apps.<cluster_name>.<base_domain>` DNS wildcard lookup. All of the application wildcard lookups must resolve to the IP address of the application ingress load balancer:

```
$ dig +noall +answer @<nameserver_ip> random.apps.<cluster_name>.<base_domain>
```

Example output

```
random.apps.ocp4.example.com. 604800 IN A 192.168.1.5
```



NOTE

In the example outputs, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

You can replace **random** with another wildcard value. For example, you can query the route to the OpenShift Container Platform console:

```
$ dig +noall +answer @<nameserver_ip> console-openshift-console.apps.<cluster_name>.<base_domain>
```

Example output

```
console-openshift-console.apps.ocp4.example.com. 604800 IN A 192.168.1.5
```

- d. Run a lookup against the bootstrap DNS record name. Check that the result points to the IP address of the bootstrap node:

```
$ dig +noall +answer @<nameserver_ip> bootstrap.<cluster_name>.<base_domain>
```

Example output

```
bootstrap.ocp4.example.com. 604800 IN A 192.168.1.96
```


- e. Use this method to perform lookups against the DNS record names for the control plane and compute nodes. Check that the results correspond to the IP addresses of each node.
2. From your installation node, run reverse DNS lookups against the IP addresses of the load balancer and the cluster nodes. Validate that the record names contained in the responses correspond to the correct components.
- a. Perform a reverse lookup against the IP address of the API load balancer. Check that the response includes the record names for the Kubernetes API and the Kubernetes internal API:

```
$ dig +noall +answer @<nameserver_ip> -x 192.168.1.5
```

Example output

```
5.1.168.192.in-addr.arpa. 604800 IN PTR api-int.ocp4.example.com. 1
5.1.168.192.in-addr.arpa. 604800 IN PTR api.ocp4.example.com. 2
```

- 1 Provides the record name for the Kubernetes internal API.
- 2 Provides the record name for the Kubernetes API.



NOTE

A PTR record is not required for the OpenShift Container Platform application wildcard. No validation step is needed for reverse DNS resolution against the IP address of the application ingress load balancer.

- b. Perform a reverse lookup against the IP address of the bootstrap node. Check that the result points to the DNS record name of the bootstrap node:

```
$ dig +noall +answer @<nameserver_ip> -x 192.168.1.96
```

Example output

```
96.1.168.192.in-addr.arpa. 604800 IN PTR bootstrap.ocp4.example.com.
```

- c. Use this method to perform reverse lookups against the IP addresses for the control plane and compute nodes. Check that the results correspond to the DNS record names of each node.

23.3.3. Installing a cluster on vSphere with user-provisioned infrastructure

In OpenShift Container Platform version 4.16, you can install a cluster on VMware vSphere infrastructure that you provision.



NOTE

OpenShift Container Platform supports deploying a cluster to a single VMware vCenter only. Deploying a cluster with machines/machine sets on multiple vCenters is not supported.



IMPORTANT

The steps for performing a user-provisioned infrastructure installation are provided as an example only. Installing a cluster with infrastructure you provide requires knowledge of the vSphere platform and the installation process of OpenShift Container Platform. Use the user-provisioned infrastructure installation instructions as a guide; you are free to create the required resources through other methods.

23.3.3.1. Prerequisites

- You have completed the tasks in [Preparing to install a cluster using user-provisioned infrastructure](#).

- You reviewed your VMware platform licenses. Red Hat does not place any restrictions on your VMware licenses, but some VMware infrastructure components require licensing.
- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You provisioned [persistent storage](#) for your cluster. To deploy a private image registry, your storage must provide **ReadWriteMany** access modes.
- Completing the installation requires that you upload the Red Hat Enterprise Linux CoreOS (RHCOS) OVA on vSphere hosts. The machine from which you complete this process requires access to port 443 on the vCenter and ESXi hosts. You verified that port 443 is accessible.
- If you use a firewall, you confirmed with the administrator that port 443 is accessible. Control plane nodes must be able to reach vCenter and ESXi hosts on port 443 for the installation to succeed.
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.

**NOTE**


Be sure to also review this site list if you are configuring a proxy.

23.3.3.2. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.

**IMPORTANT**

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

23.3.3.3. VMware vSphere region and zone enablement

You can deploy an OpenShift Container Platform cluster to multiple vSphere datacenters that run in a single VMware vCenter. Each datacenter can run multiple clusters. This configuration reduces the risk of a hardware failure or network outage that can cause your cluster to fail. To enable regions and zones, you must define multiple failure domains for your OpenShift Container Platform cluster.



IMPORTANT

The VMware vSphere region and zone enablement feature requires the vSphere Container Storage Interface (CSI) driver as the default storage driver in the cluster. As a result, the feature is only available on a newly installed cluster.

For a cluster that was upgraded from a previous release, you must enable CSI automatic migration for the cluster. You can then configure multiple regions and zones for the upgraded cluster.

The default installation configuration deploys a cluster to a single vSphere datacenter. If you want to deploy a cluster to multiple vSphere datacenters, you must create an installation configuration file that enables the region and zone feature.

The default **install-config.yaml** file includes **vcenters** and **failureDomains** fields, where you can specify multiple vSphere datacenters and clusters for your OpenShift Container Platform cluster. You can leave these fields blank if you want to install an OpenShift Container Platform cluster in a vSphere environment that consists of single datacenter.

The following list describes terms associated with defining zones and regions for your cluster:

- **Failure domain:** Establishes the relationships between a region and zone. You define a failure domain by using vCenter objects, such as a **datastore** object. A failure domain defines the vCenter location for OpenShift Container Platform cluster nodes.
- **Region:** Specifies a vCenter datacenter. You define a region by using a tag from the **openshift-region** tag category.
- **Zone:** Specifies a vCenter cluster. You define a zone by using a tag from the **openshift-zone** tag category.



NOTE

If you plan on specifying more than one failure domain in your **install-config.yaml** file, you must create tag categories, zone tags, and region tags in advance of creating the configuration file.

You must create a vCenter tag for each vCenter datacenter, which represents a region. Additionally, you must create a vCenter tag for each cluster that runs in a datacenter, which represents a zone. After you create the tags, you must attach each tag to their respective datacenters and clusters.

The following table outlines an example of the relationship among regions, zones, and tags for a configuration with multiple vSphere datacenters running in a single VMware vCenter.

Datacenter (region)	Cluster (zone)	Tags
us-east	us-east-1	us-east-1a
		us-east-1b
	us-east-2	us-east-2a
		us-east-2b

Datacenter (region)	Cluster (zone)	Tags
us-west	us-west-1	us-west-1a
		us-west-1b
	us-west-2	us-west-2a
		us-west-2b

Additional resources

- [Additional VMware vSphere configuration parameters](#)
- [Deprecated VMware vSphere configuration parameters](#)
- [vSphere automatic migration](#)
- [VMware vSphere CSI Driver Operator](#)

23.3.3.4. Manually creating the installation configuration file

Installing the cluster requires that you manually create the installation configuration file.

Prerequisites

- You have an SSH public key on your local machine to provide to the installation program. The key will be used for SSH authentication onto your cluster nodes for debugging and disaster recovery.
- You have obtained the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create an installation directory to store your required installation assets in:

```
$ mkdir <installation_directory>
```



IMPORTANT


You must create a directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

2. Customize the sample **install-config.yaml** file template that is provided and save it in the **<installation_directory>**.

**NOTE**

You must name this configuration file **install-config.yaml**.

- If you are installing a three-node cluster, modify the **install-config.yaml** file by setting the **compute.replicas** parameter to **0**. This ensures that the cluster's control planes are schedulable. For more information, see "Installing a three-node cluster on vSphere".
- Back up the **install-config.yaml** file so that you can use it to install multiple clusters.

**IMPORTANT**

The **install-config.yaml** file is consumed during the next step of the installation process. You must back it up now.

Additional resources

- [Installation configuration parameters](#)

23.3.3.4.1. Sample install-config.yaml file for VMware vSphere

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.

```

additionalTrustBundlePolicy: Proxyonly
apiVersion: v1
baseDomain: example.com 1
compute: 2
- architecture: amd64
  name: <worker_node>
  platform: {}
  replicas: 0 3
controlPlane: 4
  architecture: amd64
  name: <parent_node>
  platform: {}
  replicas: 3 5
metadata:
  creationTimestamp: null
  name: test 6
networking:
---
platform:
  vsphere:
    failureDomains: 7
    - name: <failure_domain_name>
      region: <default_region_name>
      server: <fully_qualified_domain_name>
    topology:
      computeCluster: "/<datacenter>/host/<cluster>"
      datacenter: <datacenter> 8
      datastore: "/<datacenter>/datastore/<datastore>" 9
      networks:
      - <VM_Network_name>

```



```

resourcePool: "/<datacenter>/host/<cluster>/Resources/<resourcePool>" 10
folder: "/<datacenter_name>/vm/<folder_name>/<subfolder_name>" 11
zone: <default_zone_name>
vcenters:
- datacenters:
- <datacenter>
password: <password> 12
port: 443
server: <fully_qualified_domain_name> 13
user: administrator@vsphere.local
diskType: thin 14
fips: false 15
pullSecret: '{"auths": ...}' 16
sshKey: 'ssh-ed25519 AAAA...' 17

```

- 1 The base domain of the cluster. All DNS records must be sub-domains of this base and include the cluster name.
- 2 4 The **controlPlane** section is a single mapping, but the compute section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Both sections define a single machine pool, so only one control plane is used. OpenShift Container Platform does not support defining multiple compute pools.
- 3 You must set the value of the **replicas** parameter to **0**. This parameter controls the number of workers that the cluster creates and manages for you, which are functions that the cluster does not perform when you use user-provisioned infrastructure. You must manually deploy worker machines for the cluster to use before you finish installing OpenShift Container Platform.
- 5 The number of control plane machines that you add to the cluster. Because the cluster uses this values as the number of etcd endpoints in the cluster, the value must match the number of control plane machines that you deploy.
- 6 The cluster name that you specified in your DNS records.
- 7 Establishes the relationships between a region and zone. You define a failure domain by using vCenter objects, such as a **datastore** object. A failure domain defines the vCenter location for OpenShift Container Platform cluster nodes.
- 8 The vSphere datacenter.
- 9 The path to the vSphere datastore that holds virtual machine files, templates, and ISO images.



IMPORTANT

You can specify the path of any datastore that exists in a datastore cluster. By default, Storage vMotion is automatically enabled for a datastore cluster. Red Hat does not support Storage vMotion, so you must disable Storage vMotion to avoid data loss issues for your OpenShift Container Platform cluster.

If you must specify VMs across multiple datastores, use a **datastore** object to specify a failure domain in your cluster's **install-config.yaml** configuration file. For more information, see "VMware vSphere region and zone enablement".

- 10 Optional: For installer-provisioned infrastructure, the absolute path of an existing resource pool where the installation program creates the virtual machines, for example,
- 11 Optional: For installer-provisioned infrastructure, the absolute path of an existing folder where the installation program creates the virtual machines, for example, `/<datacenter_name>/vm/<folder_name>/<subfolder_name>`. If you do not provide this value, the installation program creates a top-level folder in the datacenter virtual machine folder that is named with the infrastructure ID. If you are providing the infrastructure for the cluster and you do not want to use the default **StorageClass** object, named **thin**, you can omit the **folder** parameter from the **install-config.yaml** file.
- 12 The password associated with the vSphere user.
- 13 The fully-qualified hostname or IP address of the vCenter server.



IMPORTANT

The Cloud Controller Manager Operator performs a connectivity check on a provided hostname or IP address. Ensure that you specify a hostname or an IP address to a reachable vCenter server. If you provide metadata to a non-existent vCenter server, installation of the cluster fails at the bootstrap stage.

- 14 The vSphere disk provisioning method.
- 15 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.



IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 16 The pull secret that you obtained from [OpenShift Cluster Manager](#). This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.
- 17 The public portion of the default SSH key for the **core** user in Red Hat Enterprise Linux CoreOS (RHCOS).

23.3.3.4.2. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations. You must include vCenter's IP address and the IP range that you use for its machines.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.

- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and

**NOTE**

The installation program does not support the proxy **readinessEndpoints** field.

**NOTE**

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.

**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

23.3.3.4.3. Configuring regions and zones for a VMware vCenter

You can modify the default installation configuration file, so that you can deploy an OpenShift Container Platform cluster to multiple vSphere datacenters that run in a single VMware vCenter.

The default **install-config.yaml** file configuration from the previous release of OpenShift Container Platform is deprecated. You can continue to use the deprecated default configuration, but the **openshift-installer** will prompt you with a warning message that indicates the use of deprecated fields in the configuration file.

**IMPORTANT**

The example uses the **govc** command. The **govc** command is an open source command available from VMware; it is not available from Red Hat. The Red Hat support team does not maintain the **govc** command. Instructions for downloading and installing **govc** are found on the VMware documentation website

Prerequisites

- You have an existing **install-config.yaml** installation configuration file.



IMPORTANT

You must specify at least one failure domain for your OpenShift Container Platform cluster, so that you can provision datacenter objects for your VMware vCenter server. Consider specifying multiple failure domains if you need to provision virtual machine nodes in different datacenters, clusters, datastores, and other components. To enable regions and zones, you must define multiple failure domains for your OpenShift Container Platform cluster.

Procedure

1. Enter the following **govc** command-line tool commands to create the **openshift-region** and **openshift-zone** vCenter tag categories:



IMPORTANT

If you specify different names for the **openshift-region** and **openshift-zone** vCenter tag categories, the installation of the OpenShift Container Platform cluster fails.

```
$ govc tags.category.create -d "OpenShift region" openshift-region
```

```
$ govc tags.category.create -d "OpenShift zone" openshift-zone
```

2. To create a region tag for each region vSphere datacenter where you want to deploy your cluster, enter the following command in your terminal:

```
$ govc tags.create -c <region_tag_category> <region_tag>
```

3. To create a zone tag for each vSphere cluster where you want to deploy your cluster, enter the following command:

```
$ govc tags.create -c <zone_tag_category> <zone_tag>
```

4. Attach region tags to each vCenter datacenter object by entering the following command:

```
$ govc tags.attach -c <region_tag_category> <region_tag_1> /<datacenter_1>
```

5. Attach the zone tags to each vCenter datacenter object by entering the following command:

```
$ govc tags.attach -c <zone_tag_category> <zone_tag_1> /<datacenter_1>/host/vcs-mdcnc-workload-1
```

6. Change to the directory that contains the installation program and initialize the cluster deployment according to your chosen installation requirements.

Sample `install-config.yaml` file with multiple datacenters defined in a vSphere center

```
---
compute:
---
vsphere:
```

```

zones:
  - "<machine_pool_zone_1>"
  - "<machine_pool_zone_2>"
---
controlPlane:
---
vsphere:
  zones:
    - "<machine_pool_zone_1>"
    - "<machine_pool_zone_2>"
---
platform:
  vsphere:
    vcenters:
---
  datacenters:
    - <datacenter1_name>
    - <datacenter2_name>
  failureDomains:
  - name: <machine_pool_zone_1>
    region: <region_tag_1>
    zone: <zone_tag_1>
    server: <fully_qualified_domain_name>
    topology:
      datacenter: <datacenter1>
      computeCluster: "/<datacenter1>/host/<cluster1>"
      networks:
        - <VM_Network1_name>
      datastore: "/<datacenter1>/datastore/<datastore1>"
      resourcePool: "/<datacenter1>/host/<cluster1>/Resources/<resourcePool1>"
      folder: "/<datacenter1>/vm/<folder1>"
  - name: <machine_pool_zone_2>
    region: <region_tag_2>
    zone: <zone_tag_2>
    server: <fully_qualified_domain_name>
    topology:
      datacenter: <datacenter2>
      computeCluster: "/<datacenter2>/host/<cluster2>"
      networks:
        - <VM_Network2_name>
      datastore: "/<datacenter2>/datastore/<datastore2>"
      resourcePool: "/<datacenter2>/host/<cluster2>/Resources/<resourcePool2>"
      folder: "/<datacenter2>/vm/<folder2>"
---

```

23.3.3.5. Creating the Kubernetes manifest and Ignition config files

Because you must modify some cluster definition files and manually start the cluster machines, you must generate the Kubernetes manifest and Ignition config files that the cluster needs to configure the machines.

The installation configuration file transforms into the Kubernetes manifests. The manifests wrap into the Ignition configuration files, which are later used to configure the cluster machines.



IMPORTANT

- The Ignition config files that the OpenShift Container Platform installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

Prerequisites

- You obtained the OpenShift Container Platform installation program.
- You created the **install-config.yaml** installation configuration file.

Procedure

1. Change to the directory that contains the OpenShift Container Platform installation program and generate the Kubernetes manifests for the cluster:

```
$ ./openshift-install create manifests --dir <installation_directory> 1
```

- 1 For **<installation_directory>**, specify the installation directory that contains the **install-config.yaml** file you created.

2. Remove the Kubernetes manifest files that define the control plane machines, compute machine sets, and control plane machine sets:

```
$ rm -f openshift/99_openshift-cluster-api_master-machines-*.yaml openshift/99_openshift-cluster-api_worker-machineset-*.yaml openshift/99_openshift-machine-api_master-control-plane-machine-set.yaml
```

Because you create and manage these resources yourself, you do not have to initialize them.

- You can preserve the compute machine set files to create compute machines by using the machine API, but you must update references to them to match your environment.



WARNING

If you are installing a three-node cluster, skip the following step to allow the control plane nodes to be schedulable.



IMPORTANT

When you configure control plane nodes from the default unschedulable to schedulable, additional subscriptions are required. This is because control plane nodes then become compute nodes.

3. Check that the **mastersSchedulable** parameter in the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** Kubernetes manifest file is set to **false**. This setting prevents pods from being scheduled on the control plane machines:
 - a. Open the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** file.
 - b. Locate the **mastersSchedulable** parameter and ensure that it is set to **false**.
 - c. Save and exit the file.
4. To create the Ignition configuration files, run the following command from the directory that contains the installation program:

```
$ ./openshift-install create ignition-configs --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the same installation directory.

Ignition config files are created for the bootstrap, control plane, and compute nodes in the installation directory. The **kubeadmin-password** and **kubeconfig** files are created in the **./<installation_directory>/auth** directory:

```
.
├── auth
│   ├── kubeadmin-password
│   └── kubeconfig
├── bootstrap.ign
├── master.ign
├── metadata.json
└── worker.ign
```

23.3.3.6. Extracting the infrastructure name

The Ignition config files contain a unique cluster identifier that you can use to uniquely identify your cluster in VMware vSphere. If you plan to use the cluster identifier as the name of your virtual machine folder, you must extract it.

Prerequisites

- You obtained the OpenShift Container Platform installation program and the pull secret for your cluster.
- You generated the Ignition config files for your cluster.
- You installed the **jq** package.

Procedure

- To extract and view the infrastructure name from the Ignition config file metadata, run the following command:

```
$ jq -r .infraID <installation_directory>/metadata.json 1
```

- 1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

Example output

```
openshift-vw9j6 1
```

- 1 The output of this command is your cluster name and a random string.

23.3.3.7. Installing RHCOS and starting the OpenShift Container Platform bootstrap process

To install OpenShift Container Platform on user-provisioned infrastructure on VMware vSphere, you must install Red Hat Enterprise Linux CoreOS (RHCOS) on vSphere hosts. When you install RHCOS, you must provide the Ignition config file that was generated by the OpenShift Container Platform installation program for the type of machine you are installing. If you have configured suitable networking, DNS, and load balancing infrastructure, the OpenShift Container Platform bootstrap process begins automatically after the RHCOS machines have rebooted.

Prerequisites

- You have obtained the Ignition config files for your cluster.
- You have access to an HTTP server that you can access from your computer and that the machines that you create can access.
- You have created a [vSphere cluster](#).

Procedure

1. Upload the bootstrap Ignition config file, which is named **<installation_directory>/bootstrap.ign**, that the installation program created to your HTTP server. Note the URL of this file.
2. Save the following secondary Ignition config file for your bootstrap node to your computer as **<installation_directory>/merge-bootstrap.ign**:

```
{
  "ignition": {
    "config": {
      "merge": [
        {
          "source": "<bootstrap_ignition_config_url>", 1
          "verification": {}
        }
      ]
    }
  },
```

```

    "timeouts": {},
    "version": "3.2.0"
  },
  "networkd": {},
  "passwd": {},
  "storage": {},
  "systemd": {}
}

```

- 1 Specify the URL of the bootstrap Ignition config file that you hosted.

When you create the virtual machine (VM) for the bootstrap machine, you use this Ignition config file.

3. Locate the following Ignition config files that the installation program created:

- **<installation_directory>/master.ign**
- **<installation_directory>/worker.ign**
- **<installation_directory>/merge-bootstrap.ign**

4. Convert the Ignition config files to Base64 encoding. Later in this procedure, you must add these files to the extra configuration parameter **guestinfo.ignition.config.data** in your VM. For example, if you use a Linux operating system, you can use the **base64** command to encode the files.

```
$ base64 -w0 <installation_directory>/master.ign > <installation_directory>/master.64
```

```
$ base64 -w0 <installation_directory>/worker.ign > <installation_directory>/worker.64
```


```
$ base64 -w0 <installation_directory>/merge-bootstrap.ign > <installation_directory>/merge-bootstrap.64
```



IMPORTANT

If you plan to add more compute machines to your cluster after you finish installation, do not delete these files.

5. Obtain the RHCOS OVA image. Images are available from the [RHCOS image mirror](#) page.



IMPORTANT

The RHCOS images might not change with every release of OpenShift Container Platform. You must download an image with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Use the image version that matches your OpenShift Container Platform version if it is available.

The filename contains the OpenShift Container Platform version number in the format **rhcos-vmware.<architecture>.ova**.

6. In the vSphere Client, create a folder in your datacenter to store your VMs.

- a. Click the **VMs and Templates** view.
 - b. Right-click the name of your datacenter.
 - c. Click **New Folder → New VM and Template Folder**.
 - d. In the window that is displayed, enter the folder name. If you did not specify an existing folder in the **install-config.yaml** file, then create a folder with the same name as the infrastructure ID. You use this folder name so vCenter dynamically provisions storage in the appropriate location for its Workspace configuration.
7. In the vSphere Client, create a template for the OVA image and then clone the template as needed.



NOTE

In the following steps, you create a template and then clone the template for all of your cluster machines. You then provide the location for the Ignition config file for that cloned machine type when you provision the VMs.


- a. From the **Hosts and Clusters** tab, right-click your cluster name and select **Deploy OVF Template**.
- b. On the **Select an OVF** tab, specify the name of the RHCOS OVA file that you downloaded.
- c. On the **Select a name and folder** tab, set a **Virtual machine name** for your template, such as **Template-RHCOS**. Click the name of your vSphere cluster and select the folder you created in the previous step.
- d. On the **Select a compute resource** tab, click the name of your vSphere cluster.
- e. On the **Select storage** tab, configure the storage options for your VM.
 - Select **Thin Provision** or **Thick Provision**, based on your storage preferences.
 - Select the datastore that you specified in your **install-config.yaml** file.
 - If you want to encrypt your virtual machines, select **Encrypt this virtual machine**. See the section titled "Requirements for encrypting virtual machines" for more information.
- f. On the **Select network** tab, specify the network that you configured for the cluster, if available.
- g. When creating the OVF template, do not specify values on the **Customize template** tab or configure the template any further.



IMPORTANT

Do not start the original VM template. The VM template must remain off and must be cloned for new RHCOS machines. Starting the VM template configures the VM template as a VM on the platform, which prevents it from being used as a template that compute machine sets can apply configurations to.

8. Optional: Update the configured virtual hardware version in the VM template, if necessary. Follow [Upgrading a virtual machine to the latest hardware version](#) in the VMware documentation for more information.



IMPORTANT

It is recommended that you update the hardware version of the VM template to version 15 before creating VMs from it, if necessary. Using hardware version 13 for your cluster nodes running on vSphere is now deprecated. If your imported template defaults to hardware version 13, you must ensure that your ESXi host is on 6.7U3 or later before upgrading the VM template to hardware version 15. If your vSphere version is less than 6.7U3, you can skip this upgrade step; however, a future version of OpenShift Container Platform is scheduled to remove support for hardware version 13 and vSphere versions less than 6.7U3.

9. After the template deploys, deploy a VM for a machine in the cluster.
 - a. Right-click the template name and click **Clone → Clone to Virtual Machine**
 - b. On the **Select a name and folder** tab, specify a name for the VM. You might include the machine type in the name, such as **control-plane-0** or **compute-1**.



NOTE

Ensure that all virtual machine names across a vSphere installation are unique.

- c. On the **Select a name and folder** tab, select the name of the folder that you created for the cluster.
- d. On the **Select a compute resource** tab, select the name of a host in your datacenter.
- e. On the **Select clone options** tab, select **Customize this virtual machine's hardware**
- f. On the **Customize hardware** tab, click **Advanced Parameters**.



IMPORTANT

The following configuration suggestions are for example purposes only. As a cluster administrator, you must configure resources according to the resource demands placed on your cluster. To best manage cluster resources, consider creating a resource pool from the cluster's root resource pool.

- Optional: Override default DHCP networking in vSphere. To enable static IP networking:
 - Set your static IP configuration:

Example command

```
$ export IPCFG="ip=<ip>::<gateway>:<netmask>:<hostname>:<iface>:none
nameserver=svr1 [nameserver=svr2 [nameserver=svr3 [...]]]"
```

Example command

```
$ export IPCFG="ip=192.168.100.101::192.168.100.254:255.255.255.0::none
nameserver=8.8.8.8"
```

- Set the **guestinfo.afterburn.initrd.network-kargs** property before you boot a VM from an OVA in vSphere:

Example command

```
$ govc vm.change -vm "<vm_name>" -e "guestinfo.afterburn.initrd.network-
kargs=${IPCFG}"
```


- Add the following configuration parameter names and values by specifying data in the **Attribute** and **Values** fields. Ensure that you select the **Add** button for each parameter that you create.
 - **guestinfo.ignition.config.data**: Locate the base-64 encoded files that you created previously in this procedure, and paste the contents of the base64-encoded Ignition config file for this machine type.
 - **guestinfo.ignition.config.data.encoding**: Specify **base64**.
 - **disk.EnableUUID**: Specify **TRUE**.
 - **stealclock.enable**: If this parameter was not defined, add it and specify **TRUE**.
 - Create a child resource pool from the cluster's root resource pool. Perform resource allocation in this child resource pool.
- g. In the **Virtual Hardware** panel of the **Customize hardware** tab, modify the specified values as required. Ensure that the amount of RAM, CPU, and disk storage meets the minimum requirements for the machine type.
- h. Complete the remaining configuration steps. On clicking the **Finish** button, you have completed the cloning operation.
- i. From the **Virtual Machines** tab, right-click on your VM and then select **Power → Power On**.
- j. Check the console output to verify that Ignition ran.

Example command

```
Ignition: ran on 2022/03/14 14:48:33 UTC (this boot)
Ignition: user-provided config was applied
```

Next steps

- Create the rest of the machines for your cluster by following the preceding steps for each machine.



IMPORTANT

You must create the bootstrap and control plane machines at this time. Because some pods are deployed on compute machines by default, also create at least two compute machines before you install the cluster.

23.3.3.8. Adding more compute machines to a cluster in vSphere

You can add more compute machines to a user-provisioned OpenShift Container Platform cluster on VMware vSphere.

After your vSphere template deploys in your OpenShift Container Platform cluster, you can deploy a virtual machine (VM) for a machine in that cluster.



NOTE

If you are installing a three-node cluster, skip this step. A three-node cluster consists of three control plane machines, which also act as compute machines.

Prerequisites

- Obtain the base64-encoded Ignition file for your compute machines.
- You have access to the vSphere template that you created for your cluster.

Procedure

1. Right-click the template's name and click **Clone** → **Clone to Virtual Machine**
2. On the **Select a name and folder** tab, specify a name for the VM. You might include the machine type in the name, such as **compute-1**.



NOTE

Ensure that all virtual machine names across a vSphere installation are unique.

3. On the **Select a name and folder** tab, select the name of the folder that you created for the cluster.
4. On the **Select a compute resource** tab, select the name of a host in your datacenter.
5. On the **Select storage** tab, select storage for your configuration and disk files.
6. On the **Select clone options** tab, select **Customize this virtual machine's hardware**
7. On the **Customize hardware** tab, click **Advanced Parameters**.
 - Add the following configuration parameter names and values by specifying data in the **Attribute** and **Values** fields. Ensure that you select the **Add** button for each parameter that you create.
 - **guestinfo.ignition.config.data**: Paste the contents of the base64-encoded compute Ignition config file for this machine type.
 - **guestinfo.ignition.config.data.encoding**: Specify **base64**.
 - **disk.EnableUUID**: Specify **TRUE**.
8. In the **Virtual Hardware** panel of the **Customize hardware** tab, modify the specified values as required. Ensure that the amount of RAM, CPU, and disk storage meets the minimum requirements for the machine type. If many networks exist, select **Add New Device** > **Network**

Adapter, and then enter your network information in the fields provided by the **New Network** menu item.

9. Complete the remaining configuration steps. On clicking the **Finish** button, you have completed the cloning operation.
10. From the **Virtual Machines** tab, right-click on your VM and then select **Power → Power On**.

Next steps


- Continue to create more compute machines for your cluster.

23.3.3.9. Disk partitioning

In most cases, data partitions are originally created by installing RHCOS, rather than by installing another operating system. In such cases, the OpenShift Container Platform installer should be allowed to configure your disk partitions.


However, there are two cases where you might want to intervene to override the default partitioning when installing an OpenShift Container Platform node:

- **Create separate partitions:** For greenfield installations on an empty disk, you might want to add separate storage to a partition. This is officially supported for making **/var** or a subdirectory of **/var**, such as **/var/lib/etcd**, a separate partition, but not both.



IMPORTANT

For disk sizes larger than 100GB, and especially disk sizes larger than 1TB, create a separate **/var** partition. See "Creating a separate **/var** partition" and this [Red Hat Knowledgebase article](#) for more information.



IMPORTANT

Kubernetes supports only two file system partitions. If you add more than one partition to the original configuration, Kubernetes cannot monitor all of them.

- **Retain existing partitions:** For a brownfield installation where you are reinstalling OpenShift Container Platform on an existing node and want to retain data partitions installed from your previous operating system, there are both boot arguments and options to **coreos-installer** that allow you to retain existing data partitions.


Creating a separate **/var** partition

In general, disk partitioning for OpenShift Container Platform should be left to the installer. However, there are cases where you might want to create separate partitions in a part of the filesystem that you expect to grow.

OpenShift Container Platform supports the addition of a single partition to attach storage to either the **/var** partition or a subdirectory of **/var**. For example:

- **/var/lib/containers:** Holds container-related content that can grow as more images and containers are added to a system.
- **/var/lib/etcd:** Holds data that you might want to keep separate for purposes such as performance optimization of etcd storage.

- **/var**: Holds data that you might want to keep separate for purposes such as auditing.



IMPORTANT

For disk sizes larger than 100GB, and especially larger than 1TB, create a separate **/var** partition.

Storing the contents of a **/var** directory separately makes it easier to grow storage for those areas as needed and reinstall OpenShift Container Platform at a later date and keep that data intact. With this method, you will not have to pull all your containers again, nor will you have to copy massive log files when you update systems.

Because **/var** must be in place before a fresh installation of Red Hat Enterprise Linux CoreOS (RHCOS), the following procedure sets up the separate **/var** partition by creating a machine config manifest that is inserted during the **openshift-install** preparation phases of an OpenShift Container Platform installation.

Procedure

1. Create a directory to hold the OpenShift Container Platform installation files:

```
$ mkdir $HOME/clusterconfig
```

2. Run **openshift-install** to create a set of files in the **manifest** and **openshift** subdirectories. Answer the system questions as you are prompted:

```
$ openshift-install create manifests --dir $HOME/clusterconfig
? SSH Public Key ...
$ ls $HOME/clusterconfig/openshift/
99_kubeadmin-password-secret.yaml
99_openshift-cluster-api_master-machines-0.yaml
99_openshift-cluster-api_master-machines-1.yaml
99_openshift-cluster-api_master-machines-2.yaml
...
```

3. Create a Butane config that configures the additional partition. For example, name the file **\$HOME/clusterconfig/98-var-partition.bu**, change the disk device name to the name of the storage device on the **worker** systems, and set the storage size as appropriate. This example places the **/var** directory on a separate partition:

```
variant: openshift
version: 4.16.0
metadata:
  labels:
    machineconfiguration.openshift.io/role: worker
  name: 98-var-partition
storage:
  disks:
    - device: /dev/disk/by-id/<device_name> 1
      partitions:
        - label: var
          start_mib: <partition_start_offset> 2
          size_mib: <partition_size> 3
          number: 5
```



```

filesystems:
- device: /dev/disk/by-partlabel/var
  path: /var
  format: xfs
  mount_options: [defaults, prjquota] 4
  with_mount_unit: true

```

- 1 The storage device name of the disk that you want to partition.
- 2 When adding a data partition to the boot disk, a minimum value of 25000 mebibytes is recommended. The root file system is automatically resized to fill all available space up to the specified offset. If no value is specified, or if the specified value is smaller than the recommended minimum, the resulting root file system will be too small, and future reinstalls of RHCOS might overwrite the beginning of the data partition.
- 3 The size of the data partition in mebibytes.
- 4 The **prjquota** mount option must be enabled for filesystems used for container storage.



NOTE

When creating a separate **/var** partition, you cannot use different instance types for worker nodes, if the different instance types do not have the same device name.

4. Create a manifest from the Butane config and save it to the **clusterconfig/openshift** directory. For example, run the following command:

```
$ butane $HOME/clusterconfig/98-var-partition.bu -o $HOME/clusterconfig/openshift/98-var-partition.yaml
```

5. Run **openshift-install** again to create Ignition configs from a set of files in the **manifest** and **openshift** subdirectories:

```
$ openshift-install create ignition-configs --dir $HOME/clusterconfig
$ ls $HOME/clusterconfig/
auth bootstrap.ign master.ign metadata.json worker.ign
```

Now you can use the Ignition config files as input to the vSphere installation procedures to install Red Hat Enterprise Linux CoreOS (RHCOS) systems.

23.3.3.10. Waiting for the bootstrap process to complete

The OpenShift Container Platform bootstrap process begins after the cluster nodes first boot into the persistent RHCOS environment that has been installed to disk. The configuration information provided through the Ignition config files is used to initialize the bootstrap process and install OpenShift Container Platform on the machines. You must wait for the bootstrap process to complete.

Prerequisites

- You have created the Ignition config files for your cluster.
- You have configured suitable network, DNS and load balancing infrastructure.

- You have obtained the installation program and generated the Ignition config files for your cluster.
- You installed RHCOS on your cluster machines and provided the Ignition config files that the OpenShift Container Platform installation program generated.
- Your machines have direct internet access or have an HTTP or HTTPS proxy available.

Procedure

1. Monitor the bootstrap process:

```
$ ./openshift-install --dir <installation_directory> wait-for bootstrap-complete \ 1
--log-level=info 2
```

1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.


2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Example output

```
INFO Waiting up to 30m0s for the Kubernetes API at https://api.test.example.com:6443...
INFO API v1.29.4 up
INFO Waiting up to 30m0s for bootstrapping to complete...
INFO It is now safe to remove the bootstrap resources
```

The command succeeds when the Kubernetes API server signals that it has been bootstrapped on the control plane machines.

2. After the bootstrap process is complete, remove the bootstrap machine from the load balancer.



IMPORTANT

You must remove the bootstrap machine from the load balancer at this point. You can also remove or reformat the bootstrap machine itself.

23.3.3.11. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

23.3.3.12. Approving the certificate signing requests for your machines

When you add machines to a cluster, two pending certificate signing requests (CSRs) are generated for each machine that you added. You must confirm that these CSRs are approved or, if necessary, approve them yourself. The client requests must be approved first, followed by the server requests.

Prerequisites

- You added machines to your cluster.

Procedure

1. Confirm that the cluster recognizes the machines:

```
$ oc get nodes
```

Example output

```
NAME      STATUS  ROLES  AGE  VERSION
master-0  Ready   master 63m  v1.29.4
master-1  Ready   master 63m  v1.29.4
master-2  Ready   master 64m  v1.29.4
```

The output lists all of the machines that you created.



NOTE

The preceding output might not include the compute nodes, also known as worker nodes, until some CSRs are approved.

2. Review the pending CSRs and ensure that you see the client requests with the **Pending** or **Approved** status for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

NAME	AGE	REQUESTOR	CONDITION
csr-8b2br	15m	system:serviceaccount:openshift-machine-config-operator:node-bootstrapper	Pending
csr-8vnps	15m	system:serviceaccount:openshift-machine-config-operator:node-bootstrapper	Pending
...			

In this example, two machines are joining the cluster. You might see more approved CSRs in the list.

- If the CSRs were not approved, after all of the pending CSRs for the machines you added are in **Pending** status, approve the CSRs for your cluster machines:



NOTE

Because the CSRs rotate automatically, approve your CSRs within an hour of adding the machines to the cluster. If you do not approve them within an hour, the certificates will rotate, and more than two certificates will be present for each node. You must approve all of these certificates. After the client CSR is approved, the Kubelet creates a secondary CSR for the serving certificate, which requires manual approval. Then, subsequent serving certificate renewal requests are automatically approved by the **machine-approver** if the Kubelet requests a new certificate with identical parameters.



NOTE

For clusters running on platforms that are not machine API enabled, such as bare metal and other user-provisioned infrastructure, you must implement a method of automatically approving the kubelet serving certificate requests (CSRs). If a request is not approved, then the **oc exec**, **oc rsh**, and **oc logs** commands cannot succeed, because a serving certificate is required when the API server connects to the kubelet. Any operation that contacts the Kubelet endpoint requires this certificate approval to be in place. The method must watch for new CSRs, confirm that the CSR was submitted by the **node-bootstrapper** service account in the **system:node** or **system:admin** groups, and confirm the identity of the node.

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{"\n"}\n{{end}}{{end}}' | xargs --no-run-if-empty oc adm certificate approve
```



NOTE

Some Operators might not become available until some CSRs are approved.

4. Now that your client requests are approved, you must review the server requests for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

```
NAME      AGE   REQUESTOR                                     CONDITION
csr-bfd72 5m26s system:node:ip-10-0-50-126.us-east-2.compute.internal
Pending
csr-c57lv 5m26s system:node:ip-10-0-95-157.us-east-2.compute.internal
Pending
...
```

5. If the remaining CSRs are not approved, and are in the **Pending** status, approve the CSRs for your cluster machines:

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:


```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{"\n"}\n{{end}}{{end}}' | xargs oc adm certificate approve
```

6. After all client and server CSRs have been approved, the machines have the **Ready** status. Verify this by running the following command:

```
$ oc get nodes
```

Example output

```
NAME      STATUS   ROLES    AGE   VERSION
master-0  Ready   master   73m   v1.29.4
master-1  Ready   master   73m   v1.29.4
master-2  Ready   master   74m   v1.29.4
worker-0  Ready   worker   11m   v1.29.4
worker-1  Ready   worker   11m   v1.29.4
```



NOTE

It can take a few minutes after approval of the server CSRs for the machines to transition to the **Ready** status.

Additional information

- For more information on CSRs, see [Certificate Signing Requests](#).

23.3.3.13. Initial Operator configuration

After the control plane initializes, you must immediately configure some Operators so that they all become available.

Prerequisites

- Your control plane has initialized.

Procedure

1. Watch the cluster components come online:

```
$ watch -n5 oc get clusteroperators
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.16.0	True	False	False	19m
baremetal	4.16.0	True	False	False	37m
cloud-credential	4.16.0	True	False	False	40m
cluster-autoscaler	4.16.0	True	False	False	37m
config-operator	4.16.0	True	False	False	38m
console	4.16.0	True	False	False	26m
csi-snapshot-controller	4.16.0	True	False	False	37m
dns	4.16.0	True	False	False	37m
etcd	4.16.0	True	False	False	36m
image-registry	4.16.0	True	False	False	31m
ingress	4.16.0	True	False	False	30m
insights	4.16.0	True	False	False	31m
kube-apiserver	4.16.0	True	False	False	26m
kube-controller-manager	4.16.0	True	False	False	36m
kube-scheduler	4.16.0	True	False	False	36m
kube-storage-version-migrator	4.16.0	True	False	False	37m
machine-api	4.16.0	True	False	False	29m
machine-approver	4.16.0	True	False	False	37m
machine-config	4.16.0	True	False	False	36m
marketplace	4.16.0	True	False	False	37m
monitoring	4.16.0	True	False	False	29m
network	4.16.0	True	False	False	38m
node-tuning	4.16.0	True	False	False	37m
openshift-apiserver	4.16.0	True	False	False	32m
openshift-controller-manager	4.16.0	True	False	False	30m
openshift-samples	4.16.0	True	False	False	32m
operator-lifecycle-manager	4.16.0	True	False	False	37m
operator-lifecycle-manager-catalog	4.16.0	True	False	False	37m
operator-lifecycle-manager-packageserver	4.16.0	True	False	False	32m
service-ca	4.16.0	True	False	False	38m
storage	4.16.0	True	False	False	37m

2. Configure the Operators that are not available.

23.3.3.13.1. Image registry removed during installation

On platforms that do not provide shareable object storage, the OpenShift Image Registry Operator bootstraps itself as **Removed**. This allows **openshift-installer** to complete installations on these platform types.

After installation, you must edit the Image Registry Operator configuration to switch the **managementState** from **Removed** to **Managed**. When this has completed, you must configure storage.

23.3.3.13.2. Image registry storage configuration

The Image Registry Operator is not initially available for platforms that do not provide default storage. After installation, you must configure your registry to use storage so that the Registry Operator is made available.

Instructions are shown for configuring a persistent volume, which is required for production clusters. Where applicable, instructions are shown for configuring an empty directory as the storage location, which is available for only non-production clusters.


Additional instructions are provided for allowing the image registry to use block storage types by using the **Recreate** rollout strategy during upgrades.

23.3.3.13.2.1. Configuring registry storage for VMware vSphere

As a cluster administrator, following installation you must configure your registry to use storage.

Prerequisites


- Cluster administrator permissions.
- A cluster on VMware vSphere.
- Persistent storage provisioned for your cluster, such as Red Hat OpenShift Data Foundation.



IMPORTANT

OpenShift Container Platform supports **ReadWriteOnce** access for image registry storage when you have only one replica. **ReadWriteOnce** access also requires that the registry uses the **Recreate** rollout strategy. To deploy an image registry that supports high availability with two or more replicas, **ReadWriteMany** access is required.

- Must have "100Gi" capacity.



IMPORTANT

Testing shows issues with using the NFS server on RHEL as storage backend for core services. This includes the OpenShift Container Registry and Quay, Prometheus for monitoring storage, and Elasticsearch for logging storage. Therefore, using RHEL NFS to back PVs used by core services is not recommended.

Other NFS implementations on the marketplace might not have these issues. Contact the individual NFS implementation vendor for more information on any testing that was possibly completed against these OpenShift Container Platform core components.

Procedure

1. To configure your registry to use storage, change the **spec.storage.pvc** in the **configs.imageregistry/cluster** resource.

**NOTE**

When you use shared storage, review your security settings to prevent outside access.

2. Verify that you do not have a registry pod:

```
$ oc get pod -n openshift-image-registry -l docker-registry=default
```

Example output

```
No resources found in openshift-image-registry namespace
```

**NOTE**

If you do have a registry pod in your output, you do not need to continue with this procedure.

3. Check the registry configuration:

```
$ oc edit configs.imageregistry.operator.openshift.io
```

Example output

```
storage:
  pvc:
    claim: 1
```

- 1** Leave the **claim** field blank to allow the automatic creation of an **image-registry-storage** persistent volume claim (PVC). The PVC is generated based on the default storage class. However, be aware that the default storage class might provide ReadWriteOnce (RWO) volumes, such as a RADOS Block Device (RBD), which can cause issues when you replicate to more than one replica.

4. Check the **clusteroperator** status:

```
$ oc get clusteroperator image-registry
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED
image-registry	4.7	True	False	False

23.3.3.13.2.2. Configuring storage for the image registry in non-production clusters

You must configure storage for the Image Registry Operator. For non-production clusters, you can set the image registry to an empty directory. If you do so, all images are lost if you restart the registry.

Procedure

- To set the image registry storage to an empty directory:

```
$ oc patch configs.imageregistry.operator.openshift.io cluster --type merge --patch '{"spec": {"storage":{"emptyDir":{}}}'
```



WARNING

Configure this option for only non-production clusters.


If you run this command before the Image Registry Operator initializes its components, the **oc patch** command fails with the following error:

```
Error from server (NotFound): configs.imageregistry.operator.openshift.io "cluster" not found
```

Wait a few minutes and run the command again.

23.3.3.13.2.3. Configuring block registry storage for VMware vSphere

To allow the image registry to use block storage types such as vSphere Virtual Machine Disk (VMDK) during upgrades as a cluster administrator, you can use the **Recreate** rollout strategy.



IMPORTANT

Block storage volumes are supported but not recommended for use with image registry on production clusters. An installation where the registry is configured on block storage is not highly available because the registry cannot have more than one replica.

Procedure

- Enter the following command to set the image registry storage as a block storage type, patch the registry so that it uses the **Recreate** rollout strategy, and runs with only **1** replica:

```
$ oc patch config.imageregistry.operator.openshift.io/cluster --type=merge -p '{"spec": {"rolloutStrategy":"Recreate","replicas":1}}'
```

- Provision the PV for the block storage device, and create a PVC for that volume. The requested block volume uses the ReadWriteOnce (RWO) access mode.
 - Create a **pvc.yaml** file with the following contents to define a VMware vSphere **PersistentVolumeClaim** object:

```
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
```

```

name: image-registry-storage 1
namespace: openshift-image-registry 2
spec:
  accessModes:
  - ReadWriteOnce 3
  resources:
    requests:
      storage: 100Gi 4

```

- 1** A unique name that represents the **PersistentVolumeClaim** object.
- 2** The namespace for the **PersistentVolumeClaim** object, which is **openshift-image-registry**.
- 3** The access mode of the persistent volume claim. With **ReadWriteOnce**, the volume can be mounted with read and write permissions by a single node.
- 4** The size of the persistent volume claim.

b. Enter the following command to create the **PersistentVolumeClaim** object from the file:

```
$ oc create -f pvc.yaml -n openshift-image-registry
```

3. Enter the following command to edit the registry configuration so that it references the correct PVC:

```
$ oc edit config.imageregistry.operator.openshift.io -o yaml
```

Example output

```

storage:
  pvc:
    claim: 1

```

- 1** By creating a custom PVC, you can leave the **claim** field blank for the default automatic creation of an **image-registry-storage** PVC.

For instructions about configuring registry storage so that it references the correct PVC, see [Configuring the registry for vSphere](#).

23.3.3.14. Completing installation on user-provisioned infrastructure

After you complete the Operator configuration, you can finish installing the cluster on infrastructure that you provide.

Prerequisites

- Your control plane has initialized.
- You have completed the initial Operator configuration.

Procedure

1. Confirm that all the cluster components are online with the following command:

```
$ watch -n5 oc get clusteroperators
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.16.0	True	False	False	19m
baremetal	4.16.0	True	False	False	37m
cloud-credential	4.16.0	True	False	False	40m
cluster-autoscaler	4.16.0	True	False	False	37m
config-operator	4.16.0	True	False	False	38m
console	4.16.0	True	False	False	26m
csi-snapshot-controller	4.16.0	True	False	False	37m
dns	4.16.0	True	False	False	37m
etcd	4.16.0	True	False	False	36m
image-registry	4.16.0	True	False	False	31m
ingress	4.16.0	True	False	False	30m
insights	4.16.0	True	False	False	31m
kube-apiserver	4.16.0	True	False	False	26m
kube-controller-manager	4.16.0	True	False	False	36m
kube-scheduler	4.16.0	True	False	False	36m
kube-storage-version-migrator	4.16.0	True	False	False	37m
machine-api	4.16.0	True	False	False	29m
machine-approver	4.16.0	True	False	False	37m
machine-config	4.16.0	True	False	False	36m
marketplace	4.16.0	True	False	False	37m
monitoring	4.16.0	True	False	False	29m
network	4.16.0	True	False	False	38m
node-tuning	4.16.0	True	False	False	37m
openshift-apiserver	4.16.0	True	False	False	32m
openshift-controller-manager	4.16.0	True	False	False	30m
openshift-samples	4.16.0	True	False	False	32m
operator-lifecycle-manager	4.16.0	True	False	False	37m
operator-lifecycle-manager-catalog	4.16.0	True	False	False	37m
operator-lifecycle-manager-packageserver	4.16.0	True	False	False	32m
service-ca	4.16.0	True	False	False	38m
storage	4.16.0	True	False	False	37m

Alternatively, the following command notifies you when all of the clusters are available. It also retrieves and displays credentials:

```
$ ./openshift-install --dir <installation_directory> wait-for install-complete 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

Example output

```
INFO Waiting up to 30m0s for the cluster to initialize...
```

The command succeeds when the Cluster Version Operator finishes deploying the OpenShift Container Platform cluster from Kubernetes API server.



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

2. Confirm that the Kubernetes API server is communicating with the pods.

- To view a list of all pods, use the following command:

```
$ oc get pods --all-namespaces
```

Example output

```

NAMESPACE           NAME                                     READY  STATUS   RESTARTS  AGE
openshift-apiserver-operator  openshift-apiserver-operator-85cb746d55-zqhs8  1/1    Running   0          9m
openshift-apiserver          apiserver-67b9g                               1/1    Running   0          3m
openshift-apiserver          apiserver-ljcmx                               1/1    Running   0          1m
openshift-apiserver          apiserver-z25h4                               1/1    Running   0          2m
openshift-authentication-operator  authentication-operator-69d5d8bf84-vh2n8      1/1    Running   0          5m
...

```

- View the logs for a pod that is listed in the output of the previous command by using the following command:

```
$ oc logs <pod_name> -n <namespace> 1
```

- 1** Specify the pod name and namespace, as shown in the output of the previous command.

If the pod logs display, the Kubernetes API server can communicate with the cluster machines.

- For an installation with Fibre Channel Protocol (FCP), additional steps are required to enable multipathing. Do not enable multipathing during installation. See "Enabling multipathing with kernel arguments on RHCOS" in the *Postinstallation machine configuration tasks* documentation for more information.

You can add extra compute machines after the cluster installation is completed by following [Adding compute machines to vSphere](#).

23.3.3.15. Configuring vSphere DRS anti-affinity rules for control plane nodes

vSphere Distributed Resource Scheduler (DRS) anti-affinity rules can be configured to support higher availability of OpenShift Container Platform Control Plane nodes. Anti-affinity rules ensure that the vSphere Virtual Machines for the OpenShift Container Platform Control Plane nodes are not scheduled to the same vSphere Host.



IMPORTANT

- The following information applies to compute DRS only and does not apply to storage DRS.
- The **govc** command is an open-source command available from VMware; it is not available from Red Hat. The **govc** command is not supported by the Red Hat support.
- Instructions for downloading and installing **govc** are found on the VMware documentation website.

Create an anti-affinity rule by running the following command:

Example command

```
$ govc cluster.rule.create \
-name openshift4-control-plane-group \
-dc MyDatacenter -cluster MyCluster \
-enable \
-anti-affinity master-0 master-1 master-2
```

After creating the rule, your control plane nodes are automatically migrated by vSphere so they are not running on the same hosts. This might take some time while vSphere reconciles the new rule. Successful command completion is shown in the following procedure.



NOTE

The migration occurs automatically and might cause brief OpenShift API outage or latency until the migration finishes.

The vSphere DRS anti-affinity rules need to be updated manually in the event of a control plane VM name change or migration to a new vSphere Cluster.

Procedure

- Remove any existing DRS anti-affinity rule by running the following command:

```
$ govc cluster.rule.remove \  
-name openshift4-control-plane-group \  
-dc MyDatacenter -cluster MyCluster
```

Example Output

```
[13-10-22 09:33:24] Reconfigure /MyDatacenter/host/MyCluster...OK
```

2. Create the rule again with updated names by running the following command:

```
$ govc cluster.rule.create \  
-name openshift4-control-plane-group \  
-dc MyDatacenter -cluster MyOtherCluster \  
-enable \  
-anti-affinity master-0 master-1 master-2
```

23.3.3.16. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service

23.3.3.17. Next steps

- [Customize your cluster](#).
- If necessary, you can [opt out of remote health reporting](#).
- [Set up your registry and configure registry storage](#).
- Optional: [View the events from the vSphere Problem Detector Operator](#) to determine if the cluster has permission or storage configuration issues.
- Optional: if you created encrypted virtual machines, [create an encrypted storage class](#).

23.3.4. Installing a cluster on vSphere with network customizations

In OpenShift Container Platform version 4.16, you can install a cluster on VMware vSphere infrastructure that you provision with customized network configuration options. By customizing your network configuration, your cluster can coexist with existing IP address allocations in your environment and integrate with existing MTU and VXLAN configurations.

**NOTE**

OpenShift Container Platform supports deploying a cluster to a single VMware vCenter only. Deploying a cluster with machines/machine sets on multiple vCenters is not supported.

You must set most of the network configuration parameters during installation, and you can modify only **kubeProxy** configuration parameters in a running cluster.

**IMPORTANT**

The steps for performing a user-provisioned infrastructure installation are provided as an example only. Installing a cluster with infrastructure you provide requires knowledge of the vSphere platform and the installation process of OpenShift Container Platform. Use the user-provisioned infrastructure installation instructions as a guide; you are free to create the required resources through other methods.

23.3.4.1. Prerequisites

- You have completed the tasks in [Preparing to install a cluster using user-provisioned infrastructure](#).
- You reviewed your VMware platform licenses. Red Hat does not place any restrictions on your VMware licenses, but some VMware infrastructure components require licensing.
- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- Completing the installation requires that you upload the Red Hat Enterprise Linux CoreOS (RHCOS) OVA on vSphere hosts. The machine from which you complete this process requires access to port 443 on the vCenter and ESXi hosts. Verify that port 443 is accessible.
- If you use a firewall, you confirmed with the administrator that port 443 is accessible. Control plane nodes must be able to reach vCenter and ESXi hosts on port 443 for the installation to succeed.
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.

23.3.4.2. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.



IMPORTANT

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

23.3.4.3. VMware vSphere region and zone enablement

You can deploy an OpenShift Container Platform cluster to multiple vSphere datacenters that run in a single VMware vCenter. Each datacenter can run multiple clusters. This configuration reduces the risk of a hardware failure or network outage that can cause your cluster to fail. To enable regions and zones, you must define multiple failure domains for your OpenShift Container Platform cluster.



IMPORTANT

The VMware vSphere region and zone enablement feature requires the vSphere Container Storage Interface (CSI) driver as the default storage driver in the cluster. As a result, the feature is only available on a newly installed cluster.

For a cluster that was upgraded from a previous release, you must enable CSI automatic migration for the cluster. You can then configure multiple regions and zones for the upgraded cluster.

The default installation configuration deploys a cluster to a single vSphere datacenter. If you want to deploy a cluster to multiple vSphere datacenters, you must create an installation configuration file that enables the region and zone feature.

The default **install-config.yaml** file includes **vcenters** and **failureDomains** fields, where you can specify multiple vSphere datacenters and clusters for your OpenShift Container Platform cluster. You can leave these fields blank if you want to install an OpenShift Container Platform cluster in a vSphere environment that consists of single datacenter.

The following list describes terms associated with defining zones and regions for your cluster:

- Failure domain: Establishes the relationships between a region and zone. You define a failure domain by using vCenter objects, such as a **datastore** object. A failure domain defines the vCenter location for OpenShift Container Platform cluster nodes.
- Region: Specifies a vCenter datacenter. You define a region by using a tag from the **openshift-region** tag category.
- Zone: Specifies a vCenter cluster. You define a zone by using a tag from the **openshift-zone** tag category.



NOTE

If you plan on specifying more than one failure domain in your **install-config.yaml** file, you must create tag categories, zone tags, and region tags in advance of creating the configuration file.

You must create a vCenter tag for each vCenter datacenter, which represents a region. Additionally, you must create a vCenter tag for each cluster than runs in a datacenter, which represents a zone. After you create the tags, you must attach each tag to their respective datacenters and clusters.

The following table outlines an example of the relationship among regions, zones, and tags for a configuration with multiple vSphere datacenters running in a single VMware vCenter.

Datacenter (region)	Cluster (zone)	Tags
us-east	us-east-1	us-east-1a
		us-east-1b
	us-east-2	us-east-2a
		us-east-2b
us-west	us-west-1	us-west-1a
		us-west-1b
	us-west-2	us-west-2a
		us-west-2b

Additional resources

- [Additional VMware vSphere configuration parameters](#)
- [Deprecated VMware vSphere configuration parameters](#)
- [vSphere automatic migration](#)
- [VMware vSphere CSI Driver Operator](#)

23.3.4.4. Manually creating the installation configuration file

Installing the cluster requires that you manually create the installation configuration file.



IMPORTANT

The Cloud Controller Manager Operator performs a connectivity check on a provided hostname or IP address. Ensure that you specify a hostname or an IP address to a reachable vCenter server. If you provide metadata to a non-existent vCenter server, installation of the cluster fails at the bootstrap stage.

Prerequisites


- You have an SSH public key on your local machine to provide to the installation program. The key will be used for SSH authentication onto your cluster nodes for debugging and disaster recovery.

- You have obtained the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

1. Create an installation directory to store your required installation assets in:

```
$ mkdir <installation_directory>
```



IMPORTANT

You must create a directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

2. Customize the sample **install-config.yaml** file template that is provided and save it in the **<installation_directory>**.



NOTE

You must name this configuration file **install-config.yaml**.

3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the next step of the installation process. You must back it up now.

Additional resources

- [Installation configuration parameters](#)

23.3.4.4.1. Sample **install-config.yaml** file for VMware vSphere

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.

```
additionalTrustBundlePolicy: Proxyonly
apiVersion: v1
baseDomain: example.com 1
compute: 2
- architecture: amd64
  name: <worker_node>
  platform: {}
  replicas: 0 3
controlPlane: 4
  architecture: amd64
  name: <parent_node>
```

```

platform: {}
replicas: 3 5
metadata:
  creationTimestamp: null
  name: test 6
networking:
---
platform:
  vsphere:
    failureDomains: 7
    - name: <failure_domain_name>
      region: <default_region_name>
      server: <fully_qualified_domain_name>
    topology:
      computeCluster: "/<datacenter>/host/<cluster>"
      datacenter: <datacenter> 8
      datastore: "/<datacenter>/datastore/<datastore>" 9
      networks:
      - <VM_Network_name>
      resourcePool: "/<datacenter>/host/<cluster>/Resources/<resourcePool>" 10
      folder: "/<datacenter_name>/vm/<folder_name>/<subfolder_name>" 11
      zone: <default_zone_name>
    vcenters:
    - datacenters:
      - <datacenter>
      password: <password> 12
      port: 443
      server: <fully_qualified_domain_name> 13
      user: administrator@vsphere.local
    diskType: thin 14
  fips: false 15
  pullSecret: '{"auths": ...}' 16
  sshKey: 'ssh-ed25519 AAAA...' 17

```

- 1** The base domain of the cluster. All DNS records must be sub-domains of this base and include the cluster name.
- 2** **4** The **controlPlane** section is a single mapping, but the compute section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Both sections define a single machine pool, so only one control plane is used. OpenShift Container Platform does not support defining multiple compute pools.
- 3** You must set the value of the **replicas** parameter to **0**. This parameter controls the number of workers that the cluster creates and manages for you, which are functions that the cluster does not perform when you use user-provisioned infrastructure. You must manually deploy worker machines for the cluster to use before you finish installing OpenShift Container Platform.
- 5** The number of control plane machines that you add to the cluster. Because the cluster uses this values as the number of etcd endpoints in the cluster, the value must match the number of control plane machines that you deploy.
- 6** The cluster name that you specified in your DNS records.

- 7 Establishes the relationships between a region and zone. You define a failure domain by using vCenter objects, such as a **datastore** object. A failure domain defines the vCenter location for
- 8 The vSphere datacenter.
- 9 The path to the vSphere datastore that holds virtual machine files, templates, and ISO images.



IMPORTANT

You can specify the path of any datastore that exists in a datastore cluster. By default, Storage vMotion is automatically enabled for a datastore cluster. Red Hat does not support Storage vMotion, so you must disable Storage vMotion to avoid data loss issues for your OpenShift Container Platform cluster.

If you must specify VMs across multiple datastores, use a **datastore** object to specify a failure domain in your cluster's **install-config.yaml** configuration file. For more information, see "VMware vSphere region and zone enablement".

- 10 Optional: For installer-provisioned infrastructure, the absolute path of an existing resource pool where the installation program creates the virtual machines, for example, `/<datacenter_name>/host/<cluster_name>/Resources/<resource_pool_name>/<optional_nested_resource_pool_name>`. If you do not specify a value, resources are installed in the root of the cluster `/example_datacenter/host/example_cluster/Resources`.
- 11 Optional: For installer-provisioned infrastructure, the absolute path of an existing folder where the installation program creates the virtual machines, for example, `/<datacenter_name>/vm/<folder_name>/<subfolder_name>`. If you do not provide this value, the installation program creates a top-level folder in the datacenter virtual machine folder that is named with the infrastructure ID. If you are providing the infrastructure for the cluster and you do not want to use the default **StorageClass** object, named **thin**, you can omit the **folder** parameter from the **install-config.yaml** file.
- 12 The password associated with the vSphere user.
- 13 The fully-qualified hostname or IP address of the vCenter server.



IMPORTANT

The Cloud Controller Manager Operator performs a connectivity check on a provided hostname or IP address. Ensure that you specify a hostname or an IP address to a reachable vCenter server. If you provide metadata to a non-existent vCenter server, installation of the cluster fails at the bootstrap stage.

- 14 The vSphere disk provisioning method.
- 15 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.



IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 16 The pull secret that you obtained from [OpenShift Cluster Manager](#). This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.
- 17 The public portion of the default SSH key for the **core** user in Red Hat Enterprise Linux CoreOS (RHCOS).

23.3.4.4.2. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
```

```


httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
noProxy: example.com 3
additionalTrustBundle: | 4
  -----BEGIN CERTIFICATE-----
  <MY_TRUSTED_CA_CERT>
  -----END CERTIFICATE-----
additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5

```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations. You must include vCenter's IP address and the IP range that you use for its machines.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.

**NOTE**

The installation program does not support the proxy **readinessEndpoints** field.

**NOTE**

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.


**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

23.3.4.4.3. Configuring regions and zones for a VMware vCenter

You can modify the default installation configuration file, so that you can deploy an OpenShift Container Platform cluster to multiple vSphere datacenters that run in a single VMware vCenter.


The default **install-config.yaml** file configuration from the previous release of OpenShift Container Platform is deprecated. You can continue to use the deprecated default configuration, but the **openshift-installer** will prompt you with a warning message that indicates the use of deprecated fields in the configuration file.

**IMPORTANT**

The example uses the **govc** command. The **govc** command is an open source command available from VMware; it is not available from Red Hat. The Red Hat support team does not maintain the **govc** command. Instructions for downloading and installing **govc** are found on the VMware documentation website

Prerequisites

- You have an existing **install-config.yaml** installation configuration file.

**IMPORTANT**

You must specify at least one failure domain for your OpenShift Container Platform cluster, so that you can provision datacenter objects for your VMware vCenter server. Consider specifying multiple failure domains if you need to provision virtual machine nodes in different datacenters, clusters, datastores, and other components. To enable regions and zones, you must define multiple failure domains for your OpenShift Container Platform cluster.

Procedure

1. Enter the following **govc** command-line tool commands to create the **openshift-region** and **openshift-zone** vCenter tag categories:

**IMPORTANT**

If you specify different names for the **openshift-region** and **openshift-zone** vCenter tag categories, the installation of the OpenShift Container Platform cluster fails.

```
$ govc tags.category.create -d "OpenShift region" openshift-region
```

```
$ govc tags.category.create -d "OpenShift zone" openshift-zone
```

2. To create a region tag for each region vSphere datacenter where you want to deploy your cluster, enter the following command in your terminal:

```
$ govc tags.create -c <region_tag_category> <region_tag>
```

- To create a zone tag for each vSphere cluster where you want to deploy your cluster, enter the following command:

```
$ govc tags.create -c <zone_tag_category> <zone_tag>
```

- Attach region tags to each vCenter datacenter object by entering the following command:

```
$ govc tags.attach -c <region_tag_category> <region_tag_1> /<datacenter_1>
```

- Attach the zone tags to each vCenter datacenter object by entering the following command:

```
$ govc tags.attach -c <zone_tag_category> <zone_tag_1> /<datacenter_1>/host/vcs-mdcnc-workload-1
```

- Change to the directory that contains the installation program and initialize the cluster deployment according to your chosen installation requirements.

Sample install-config.yaml file with multiple datacenters defined in a vSphere center

```
---
compute:
---
vsphere:
  zones:
    - "<machine_pool_zone_1>"
    - "<machine_pool_zone_2>"
---
controlPlane:
---
vsphere:
  zones:
    - "<machine_pool_zone_1>"
    - "<machine_pool_zone_2>"
---
platform:
  vsphere:
    vcenters:
---
  datacenters:
    - <datacenter1_name>
    - <datacenter2_name>
  failureDomains:
    - name: <machine_pool_zone_1>
      region: <region_tag_1>
      zone: <zone_tag_1>
      server: <fully_qualified_domain_name>
  topology:
    datacenter: <datacenter1>
    computeCluster: "/<datacenter1>/host/<cluster1>"
    networks:
      - <VM_Network1_name>
    datastore: "/<datacenter1>/datastore/<datastore1>"
```



```

resourcePool: "/<datacenter1>/host/<cluster1>/Resources/<resourcePool1>"
folder: "/<datacenter1>/vm/<folder1>"
- name: <machine_pool_zone_2>
  region: <region_tag_2>
  zone: <zone_tag_2>
  server: <fully_qualified_domain_name>
  topology:
    datacenter: <datacenter2>
    computeCluster: "/<datacenter2>/host/<cluster2>"
    networks:
      - <VM_Network2_name>
    datastore: "/<datacenter2>/datastore/<datastore2>"
    resourcePool: "/<datacenter2>/host/<cluster2>/Resources/<resourcePool2>"
    folder: "/<datacenter2>/vm/<folder2>"
---
```

23.3.4.5. Network configuration phases

There are two phases prior to OpenShift Container Platform installation where you can customize the network configuration.

Phase 1

You can customize the following network-related fields in the **install-config.yaml** file before you create the manifest files:

- **networking.networkType**
- **networking.clusterNetwork**
- **networking.serviceNetwork**
- **networking.machineNetwork**

For more information on these fields, refer to *Installation configuration parameters*.



NOTE

Set the **networking.machineNetwork** to match the CIDR that the preferred NIC resides in.



IMPORTANT

The CIDR range **172.17.0.0/16** is reserved by libVirt. You cannot use this range or any range that overlaps with this range for any networks in your cluster.

Phase 2

After creating the manifest files by running **openshift-install create manifests**, you can define a customized Cluster Network Operator manifest with only the fields you want to modify. You can use the manifest to specify advanced network configuration.

You cannot override the values specified in phase 1 in the **install-config.yaml** file during phase 2. However, you can further customize the network plugin during phase 2.

23.3.4.6. Specifying advanced network configuration

You can use advanced network configuration for your network plugin to integrate your cluster into your existing network environment. You can specify advanced network configuration only before you install the cluster.



IMPORTANT

Customizing your network configuration by modifying the OpenShift Container Platform manifest files created by the installation program is not supported. Applying a manifest file that you create, as in the following procedure, is supported.

Prerequisites

- You have created the **install-config.yaml** file and completed any modifications to it.

Procedure

- Change to the directory that contains the installation program and create the manifests:

```
$ ./openshift-install create manifests --dir <installation_directory> 1
```

- 1** **<installation_directory>** specifies the name of the directory that contains the **install-config.yaml** file for your cluster.

- Create a stub manifest file for the advanced network configuration that is named **cluster-network-03-config.yml** in the **<installation_directory>/manifests/** directory:

```
apiVersion: operator.openshift.io/v1
kind: Network
metadata:
  name: cluster
spec:
```

- Specify the advanced network configuration for your cluster in the **cluster-network-03-config.yml** file, such as in the following example:

Enable IPsec for the OVN-Kubernetes network provider

```
apiVersion: operator.openshift.io/v1
kind: Network
metadata:
  name: cluster
spec:
  defaultNetwork:
    ovnKubernetesConfig:
      ipsecConfig:
        mode: Full
```

- Optional: Back up the **manifests/cluster-network-03-config.yml** file. The installation program consumes the **manifests/** directory when you create the Ignition config files.

- Remove the Kubernetes manifest files that define the control plane machines and compute machineSets:

```
$ rm -f openshift/99_openshift-cluster-api_master-machines-*.yaml openshift/99_openshift-cluster-api_worker-machineset-*.yaml
```

Because you create and manage these resources yourself, you do not have to initialize them.

- You can preserve the MachineSet files to create compute machines by using the machine API, but you must update references to them to match your environment.

23.3.4.7. Cluster Network Operator configuration

The configuration for the cluster network is specified as part of the Cluster Network Operator (CNO) configuration and stored in a custom resource (CR) object that is named **cluster**. The CR specifies the fields for the **Network** API in the **operator.openshift.io** API group.

The CNO configuration inherits the following fields during cluster installation from the **Network** API in the **Network.config.openshift.io** API group:

clusterNetwork

IP address pools from which pod IP addresses are allocated.

serviceNetwork

IP address pool for services.

defaultNetwork.type

Cluster network plugin. **OVNKubernetes** is the only supported plugin during installation.

You can specify the cluster network plugin configuration for your cluster by setting the fields for the **defaultNetwork** object in the CNO object named **cluster**.

23.3.4.7.1. Cluster Network Operator configuration object

The fields for the Cluster Network Operator (CNO) are described in the following table:

Table 23.29. Cluster Network Operator configuration object


Field	Type	Description
metadata.name	string	The name of the CNO object. This name is always cluster .
spec.clusterNetwork	array	A list specifying the blocks of IP addresses from which pod IP addresses are allocated and the subnet prefix length assigned to each individual node in the cluster. For example: <pre>spec: clusterNetwork: - cidr: 10.128.0.0/19 hostPrefix: 23 - cidr: 10.128.32.0/19 hostPrefix: 23</pre>

Field	Type	Description
spec.serviceNetwork	array	<p>A block of IP addresses for services. The OpenShift SDN and OVN-Kubernetes network plugins support only a single IP address block for the service network. For example:</p> <pre>spec: serviceNetwork: - 172.30.0.0/14</pre> <p>You can customize this field only in the install-config.yaml file before you create the manifests. The value is read-only in the manifest file.</p>
spec.defaultNetwork	object	Configures the network plugin for the cluster network.
spec.kubeProxyConfig	object	The fields for this object specify the kube-proxy configuration. If you are using the OVN-Kubernetes cluster network plugin, the kube-proxy configuration has no effect.

defaultNetwork object configuration

The values for the **defaultNetwork** object are defined in the following table:

Table 23.30. **defaultNetwork** object

Field	Type	Description
type	string	<p>OVNKubernetes. The Red Hat OpenShift Networking network plugin is selected during installation. This value cannot be changed after cluster installation.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p>NOTE</p> <p>OpenShift Container Platform uses the OVN-Kubernetes network plugin by default. OpenShift SDN is no longer available as an installation choice for new clusters.</p> </div> </div>
ovnKubernetesConfig	object	This object is only valid for the OVN-Kubernetes network plugin.

Configuration for the OVN-Kubernetes network plugin

The following table describes the configuration fields for the OVN-Kubernetes network plugin:

Table 23.31. **ovnKubernetesConfig** object


Field	Type	Description
mtu	integer	<p>The maximum transmission unit (MTU) for the Geneve (Generic Network Virtualization Encapsulation) overlay network. This is detected automatically based on the MTU of the primary network interface. You do not normally need to override the detected MTU.</p> <p>If the auto-detected value is not what you expect it to be, confirm that the MTU on the primary network interface on your nodes is correct. You cannot use this option to change the MTU value of the primary network interface on the nodes.</p> <p>If your cluster requires different MTU values for different nodes, you must set this value to 100 less than the lowest MTU value in your cluster. For example, if some nodes in your cluster have an MTU of 9001, and some have an MTU of 1500, you must set this value to 1400.</p>
genevePort	integer	The port to use for all Geneve packets. The default value is 6081 . This value cannot be changed after cluster installation.
ipsecConfig	object	Specify a configuration object for customizing the IPsec configuration.
ipv4	object	Specifies a configuration object for IPv4 settings.
ipv6	object	Specifies a configuration object for IPv6 settings.
policyAuditConfig	object	Specify a configuration object for customizing network policy audit logging. If unset, the defaults audit log settings are used.
gatewayConfig	object	<p>Optional: Specify a configuration object for customizing how egress traffic is sent to the node gateway.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p>NOTE</p> <p>While migrating egress traffic, you can expect some disruption to workloads and service traffic until the Cluster Network Operator (CNO) successfully rolls out the changes.</p> </div> </div>

Table 23.32. `ovnKubernetesConfig.ipv4` object

Field	Type	Description
-------	------	-------------

Field	Type	Description
internalTransitSwitchSubnet	string	<p>If your existing network infrastructure overlaps with the 100.88.0.0/16 IPv4 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. The subnet for the distributed transit switch that enables east-west traffic. This subnet cannot overlap with any other subnets used by OVN-Kubernetes or on the host itself. It must be large enough to accommodate one IP address per node in your cluster.</p> <p>The default value is 100.88.0.0/16.</p>
internalJoinSubnet	string	<p>If your existing network infrastructure overlaps with the 100.64.0.0/16 IPv4 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. You must ensure that the IP address range does not overlap with any other subnet used by your OpenShift Container Platform installation. The IP address range must be larger than the maximum number of nodes that can be added to the cluster. For example, if the clusterNetwork.cidr value is 10.128.0.0/14 and the clusterNetwork.hostPrefix value is /23, then the maximum number of nodes is $2^{(23-14)}=512$.</p> <p>The default value is 100.64.0.0/16.</p>

Table 23.33. `ovnKubernetesConfig.ipv6` object

Field	Type	Description
internalTransitSwitchSubnet	string	<p>If your existing network infrastructure overlaps with the fd97::<!--64</b--> IPv6 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. The subnet for the distributed transit switch that enables east-west traffic. This subnet cannot overlap with any other subnets used by OVN-Kubernetes or on the host itself. It must be large enough to accommodate one IP address per node in your cluster.</p> <p>The default value is fd97::<!--64</b-->.</p>
internalJoinSubnet	string	<p>If your existing network infrastructure overlaps with the fd98::<!--64</b--> IPv6 subnet, you can specify a different IP address range for internal use by OVN-Kubernetes. You must ensure that the IP address range does not overlap with any other subnet used by your OpenShift Container Platform installation. The IP address range must be larger than the maximum number of nodes that can be added to the cluster.</p> <p>The default value is fd98::<!--64</b-->.</p>

Table 23.34. `policyAuditConfig` object

Field	Type	Description
rateLimit	integer	The maximum number of messages to generate every second per node. The default value is 20 messages per second.
maxFileSize	integer	The maximum size for the audit log in bytes. The default value is 50000000 or 50 MB.
maxLogFiles	integer	The maximum number of log files that are retained.
destination	string	One of the following additional audit log targets: libc The libc syslog() function of the journald process on the host. udp:<host>:<port> A syslog server. Replace <host>:<port> with the host and port of the syslog server. unix:<file> A Unix Domain Socket file specified by <file> . null Do not send the audit logs to any additional target.
syslogFacility	string	The syslog facility, such as kern , as defined by RFC5424. The default value is local0 .

Table 23.35. gatewayConfig object

Field	Type	Description
routingViaHost	boolean	Set this field to true to send egress traffic from pods to the host networking stack. For highly-specialized installations and applications that rely on manually configured routes in the kernel routing table, you might want to route egress traffic to the host networking stack. By default, egress traffic is processed in OVN to exit the cluster and is not affected by specialized routes in the kernel routing table. The default value is false . This field has an interaction with the Open vSwitch hardware offloading feature. If you set this field to true , you do not receive the performance benefits of the offloading because egress traffic is processed by the host networking stack.
ipForwarding	object	You can control IP forwarding for all traffic on OVN-Kubernetes managed interfaces by using the ipForwarding specification in the Network resource. Specify Restricted to only allow IP forwarding for Kubernetes related traffic. Specify Global to allow forwarding of all IP traffic. For new installations, the default is Restricted . For updates to OpenShift Container Platform 4.14 or later, the default is Global .

Field	Type	Description
ipv4	object	Optional: Specify an object to configure the internal OVN-Kubernetes masquerade address for host to service traffic for IPv4 addresses.
ipv6	object	Optional: Specify an object to configure the internal OVN-Kubernetes masquerade address for host to service traffic for IPv6 addresses.

Table 23.36. gatewayConfig.ipv4 object

Field	Type	Description
internalMasqueradeSubnet	string	The masquerade IPv4 addresses that are used internally to enable host to service traffic. The host is configured with these IP addresses as well as the shared gateway bridge interface. The default value is 169.254.169.0/29 .

Table 23.37. gatewayConfig.ipv6 object

Field	Type	Description
internalMasqueradeSubnet	string	The masquerade IPv6 addresses that are used internally to enable host to service traffic. The host is configured with these IP addresses as well as the shared gateway bridge interface. The default value is fd69::/125 .

Table 23.38. ipsecConfig object

Field	Type	Description
mode	string	Specifies the behavior of the IPsec implementation. Must be one of the following values: <ul style="list-style-type: none"> ● Disabled: IPsec is not enabled on cluster nodes. ● External: IPsec is enabled for network traffic with external hosts. ● Full: IPsec is enabled for pod traffic and network traffic with external hosts.

Example OVN-Kubernetes configuration with IPsec enabled

```
defaultNetwork:
  type: OVNKubernetes
  ovnKubernetesConfig:
    mtu: 1400
```



```
genevePort: 6081
ipsecConfig:
mode: Full
```




IMPORTANT

Using OVNKubernetes can lead to a stack exhaustion problem on IBM Power®.

kubeProxyConfig object configuration (OpenShiftSDN container network interface only)

The values for the **kubeProxyConfig** object are defined in the following table:

Table 23.39. kubeProxyConfig object

Field	Type	Description
iptablesSyncPeriod	string	<p>The refresh period for iptables rules. The default value is 30s. Valid suffixes include s, m, and h and are described in the Go time package documentation.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p>NOTE</p> <p>Because of performance improvements introduced in OpenShift Container Platform 4.3 and greater, adjusting the iptablesSyncPeriod parameter is no longer necessary.</p> </div> </div>
proxyArguments.iptables-min-sync-period	array	<p>The minimum duration before refreshing iptables rules. This field ensures that the refresh does not happen too frequently. Valid suffixes include s, m, and h and are described in the Go time package. The default value is:</p> <pre>kubeProxyConfig: proxyArguments: iptables-min-sync-period: - 0s</pre>

23.3.4.8. Creating the Ignition config files

Because you must manually start the cluster machines, you must generate the Ignition config files that the cluster needs to make its machines.



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

Prerequisites

- Obtain the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure

- Obtain the Ignition config files:

```
$ ./openshift-install create ignition-configs --dir <installation_directory> 1
```

- 1 For **<installation_directory>**, specify the directory name to store the files that the installation program creates.



IMPORTANT

If you created an **install-config.yaml** file, specify the directory that contains it. Otherwise, specify an empty directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

The following files are generated in the directory:

```
.
├── auth
│   ├── kubeadmin-password
│   └── kubeconfig
├── bootstrap.ign
├── master.ign
├── metadata.json
└── worker.ign
```

23.3.4.9. Extracting the infrastructure name

The Ignition config files contain a unique cluster identifier that you can use to uniquely identify your cluster in VMware vSphere. If you plan to use the cluster identifier as the name of your virtual machine folder, you must extract it.

Prerequisites

- You obtained the OpenShift Container Platform installation program and the pull secret for your cluster.
- You generated the Ignition config files for your cluster.
- You installed the **jq** package.

Procedure

- To extract and view the infrastructure name from the Ignition config file metadata, run the following command:

```
$ jq -r .infraID <installation_directory>/metadata.json 1
```

- 1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

Example output

```
openshift-vw9j6 1
```

- 1 The output of this command is your cluster name and a random string.

23.3.4.10. Installing RHCOS and starting the OpenShift Container Platform bootstrap process

To install OpenShift Container Platform on user-provisioned infrastructure on VMware vSphere, you must install Red Hat Enterprise Linux CoreOS (RHCOS) on vSphere hosts. When you install RHCOS, you must provide the Ignition config file that was generated by the OpenShift Container Platform installation program for the type of machine you are installing. If you have configured suitable networking, DNS, and load balancing infrastructure, the OpenShift Container Platform bootstrap process begins automatically after the RHCOS machines have rebooted.

Prerequisites

- You have obtained the Ignition config files for your cluster.
- You have access to an HTTP server that you can access from your computer and that the machines that you create can access.
- You have created a [vSphere cluster](#).

Procedure

1. Upload the bootstrap Ignition config file, which is named **<installation_directory>/bootstrap.ign**, that the installation program created to your HTTP server. Note the URL of this file.
2. Save the following secondary Ignition config file for your bootstrap node to your computer as **<installation_directory>/merge-bootstrap.ign**:

```
{
  "ignition": {
    "config": {
      "merge": [
        {
          "source": "<bootstrap_ignition_config_url>", 1
          "verification": {}
        }
      ]
    },
    "timeouts": {},
    "version": "3.2.0"
  },
  "networkd": {},
  "passwd": {},
  "storage": {},
  "systemd": {}
}
```

- 1 Specify the URL of the bootstrap Ignition config file that you hosted.


When you create the virtual machine (VM) for the bootstrap machine, you use this Ignition config file.

3. Locate the following Ignition config files that the installation program created:
 - **<installation_directory>/master.ign**
 - **<installation_directory>/worker.ign**
 - **<installation_directory>/merge-bootstrap.ign**
4. Convert the Ignition config files to Base64 encoding. Later in this procedure, you must add these files to the extra configuration parameter **guestinfo.ignition.config.data** in your VM. For example, if you use a Linux operating system, you can use the **base64** command to encode the files.

```
$ base64 -w0 <installation_directory>/master.ign > <installation_directory>/master.64
```

```
$ base64 -w0 <installation_directory>/worker.ign > <installation_directory>/worker.64
```

```
$ base64 -w0 <installation_directory>/merge-bootstrap.ign > <installation_directory>/merge-bootstrap.64
```

**IMPORTANT**

If you plan to add more compute machines to your cluster after you finish installation, do not delete these files.

5. Obtain the RHCOS OVA image. Images are available from the [RHCOS image mirror](#) page.

**IMPORTANT**

The RHCOS images might not change with every release of OpenShift Container Platform. You must download an image with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Use the image version that matches your OpenShift Container Platform version if it is available.

The filename contains the OpenShift Container Platform version number in the format **rhcos-vmware.<architecture>.ova**.

6. In the vSphere Client, create a folder in your datacenter to store your VMs.
 - a. Click the **VMs and Templates** view.
 - b. Right-click the name of your datacenter.
 - c. Click **New Folder → New VM and Template Folder**.
 - d. In the window that is displayed, enter the folder name. If you did not specify an existing folder in the **install-config.yaml** file, then create a folder with the same name as the infrastructure ID. You use this folder name so vCenter dynamically provisions storage in the appropriate location for its Workspace configuration.
7. In the vSphere Client, create a template for the OVA image and then clone the template as needed.

**NOTE**

In the following steps, you create a template and then clone the template for all of your cluster machines. You then provide the location for the Ignition config file for that cloned machine type when you provision the VMs.

- a. From the **Hosts and Clusters** tab, right-click your cluster name and select **Deploy OVF Template**.
- b. On the **Select an OVF** tab, specify the name of the RHCOS OVA file that you downloaded.
- c. On the **Select a name and folder** tab, set a **Virtual machine name** for your template, such as **Template-RHCOS**. Click the name of your vSphere cluster and select the folder you created in the previous step.
- d. On the **Select a compute resource** tab, click the name of your vSphere cluster.
- e. On the **Select storage** tab, configure the storage options for your VM.
 - Select **Thin Provision** or **Thick Provision**, based on your storage preferences.

- Select the datastore that you specified in your **install-config.yaml** file.
 - If you want to encrypt your virtual machines, select **Encrypt this virtual machine**. See the section titled "Requirements for encrypting virtual machines" for more information.
- f. On the **Select network** tab, specify the network that you configured for the cluster, if available.
 - g. When creating the OVF template, do not specify values on the **Customize template** tab or configure the template any further.



IMPORTANT

Do not start the original VM template. The VM template must remain off and must be cloned for new RHCOS machines. Starting the VM template configures the VM template as a VM on the platform, which prevents it from being used as a template that compute machine sets can apply configurations to.

8. Optional: Update the configured virtual hardware version in the VM template, if necessary. Follow [Upgrading a virtual machine to the latest hardware version](#) in the VMware documentation for more information.



IMPORTANT

It is recommended that you update the hardware version of the VM template to version 15 before creating VMs from it, if necessary. Using hardware version 13 for your cluster nodes running on vSphere is now deprecated. If your imported template defaults to hardware version 13, you must ensure that your ESXi host is on 6.7U3 or later before upgrading the VM template to hardware version 15. If your vSphere version is less than 6.7U3, you can skip this upgrade step; however, a future version of OpenShift Container Platform is scheduled to remove support for hardware version 13 and vSphere versions less than 6.7U3.

9. After the template deploys, deploy a VM for a machine in the cluster.
 - a. Right-click the template name and click **Clone → Clone to Virtual Machine**
 - b. On the **Select a name and folder** tab, specify a name for the VM. You might include the machine type in the name, such as **control-plane-0** or **compute-1**.



NOTE

Ensure that all virtual machine names across a vSphere installation are unique.

- c. On the **Select a name and folder** tab, select the name of the folder that you created for the cluster.
- d. On the **Select a compute resource** tab, select the name of a host in your datacenter.
- e. On the **Select clone options** tab, select **Customize this virtual machine's hardware**
- f. On the **Customize hardware** tab, click **Advanced Parameters**.



IMPORTANT

The following configuration suggestions are for example purposes only. As a cluster administrator, you must configure resources according to the resource demands placed on your cluster. To best manage cluster resources, consider creating a resource pool from the cluster's root resource pool.

- Optional: Override default DHCP networking in vSphere. To enable static IP networking:
 - Set your static IP configuration:

Example command

```
$ export IPCFG="ip=<ip>::<gateway>:<netmask>:<hostname>:<iface>:none
nameserver=svr1 [nameserver=svr2 [nameserver=svr3 [...]]]"
```

Example command

```
$ export IPCFG="ip=192.168.100.101::192.168.100.254:255.255.255.0::none
nameserver=8.8.8.8"
```

- Set the **guestinfo.afterburn.initrd.network-kargs** property before you boot a VM from an OVA in vSphere:

Example command

```
$ govc vm.change -vm "<vm_name>" -e "guestinfo.afterburn.initrd.network-
kargs=${IPCFG}"
```

- Add the following configuration parameter names and values by specifying data in the **Attribute** and **Values** fields. Ensure that you select the **Add** button for each parameter that you create.
 - **guestinfo.ignition.config.data**: Locate the base-64 encoded files that you created previously in this procedure, and paste the contents of the base64-encoded Ignition config file for this machine type.
 - **guestinfo.ignition.config.data.encoding**: Specify **base64**.
 - **disk.EnableUUID**: Specify **TRUE**.
 - **stealclock.enable**: If this parameter was not defined, add it and specify **TRUE**.
 - Create a child resource pool from the cluster's root resource pool. Perform resource allocation in this child resource pool.
- g. In the **Virtual Hardware** panel of the **Customize hardware** tab, modify the specified values as required. Ensure that the amount of RAM, CPU, and disk storage meets the minimum requirements for the machine type.
- h. Complete the remaining configuration steps. On clicking the **Finish** button, you have completed the cloning operation.
- i. From the **Virtual Machines** tab, right-click on your VM and then select **Power** → **Power On**.


- j. Check the console output to verify that Ignition ran.

Example command

```
Ignition: ran on 2022/03/14 14:48:33 UTC (this boot)
Ignition: user-provided config was applied
```

Next steps

- Create the rest of the machines for your cluster by following the preceding steps for each machine.



IMPORTANT

You must create the bootstrap and control plane machines at this time. Because some pods are deployed on compute machines by default, also create at least two compute machines before you install the cluster.

23.3.4.11. Adding more compute machines to a cluster in vSphere

You can add more compute machines to a user-provisioned OpenShift Container Platform cluster on VMware vSphere.

After your vSphere template deploys in your OpenShift Container Platform cluster, you can deploy a virtual machine (VM) for a machine in that cluster.

Prerequisites

- Obtain the base64-encoded Ignition file for your compute machines.
- You have access to the vSphere template that you created for your cluster.

Procedure

1. Right-click the template's name and click **Clone** → **Clone to Virtual Machine**
2. On the **Select a name and folder** tab, specify a name for the VM. You might include the machine type in the name, such as **compute-1**.



NOTE

Ensure that all virtual machine names across a vSphere installation are unique.

3. On the **Select a name and folder** tab, select the name of the folder that you created for the cluster.
4. On the **Select a compute resource** tab, select the name of a host in your datacenter.
5. On the **Select storage** tab, select storage for your configuration and disk files.
6. On the **Select clone options** tab, select **Customize this virtual machine's hardware**
7. On the **Customize hardware** tab, click **Advanced Parameters**.

- Add the following configuration parameter names and values by specifying data in the **Attribute** and **Values** fields. Ensure that you select the **Add** button for each parameter that you create.
 - **guestinfo.ignition.config.data**: Paste the contents of the base64-encoded compute Ignition config file for this machine type.
 - **guestinfo.ignition.config.data.encoding**: Specify **base64**.
 - **disk.EnableUUID**: Specify **TRUE**.
- 8. In the **Virtual Hardware** panel of the **Customize hardware** tab, modify the specified values as required. Ensure that the amount of RAM, CPU, and disk storage meets the minimum requirements for the machine type. If many networks exist, select **Add New Device > Network Adapter**, and then enter your network information in the fields provided by the **New Network** menu item.
- 9. Complete the remaining configuration steps. On clicking the **Finish** button, you have completed the cloning operation.
- 10. From the **Virtual Machines** tab, right-click on your VM and then select **Power → Power On**.

Next steps

- Continue to create more compute machines for your cluster.

23.3.4.12. Disk partitioning

In most cases, data partitions are originally created by installing RHCOS, rather than by installing another operating system. In such cases, the OpenShift Container Platform installer should be allowed to configure your disk partitions.


However, there are two cases where you might want to intervene to override the default partitioning when installing an OpenShift Container Platform node:

- **Create separate partitions:** For greenfield installations on an empty disk, you might want to add separate storage to a partition. This is officially supported for making **/var** or a subdirectory of **/var**, such as **/var/lib/etcd**, a separate partition, but not both.



IMPORTANT

For disk sizes larger than 100GB, and especially disk sizes larger than 1TB, create a separate **/var** partition. See "Creating a separate **/var** partition" and this [Red Hat Knowledgebase article](#) for more information.



IMPORTANT

Kubernetes supports only two file system partitions. If you add more than one partition to the original configuration, Kubernetes cannot monitor all of them.

- **Retain existing partitions:** For a brownfield installation where you are reinstalling OpenShift Container Platform on an existing node and want to retain data partitions installed from your previous operating system, there are both boot arguments and options to **coreos-installer** that allow you to retain existing data partitions.

Creating a separate `/var` partition

In general, disk partitioning for OpenShift Container Platform should be left to the installer. However, there are cases where you might want to create separate partitions in a part of the filesystem that you expect to grow.

OpenShift Container Platform supports the addition of a single partition to attach storage to either the `/var` partition or a subdirectory of `/var`. For example:

- `/var/lib/containers`: Holds container-related content that can grow as more images and containers are added to a system.
- `/var/lib/etcd`: Holds data that you might want to keep separate for purposes such as performance optimization of etcd storage.
- `/var`: Holds data that you might want to keep separate for purposes such as auditing.



IMPORTANT

For disk sizes larger than 100GB, and especially larger than 1TB, create a separate `/var` partition.

Storing the contents of a `/var` directory separately makes it easier to grow storage for those areas as needed and reinstall OpenShift Container Platform at a later date and keep that data intact. With this method, you will not have to pull all your containers again, nor will you have to copy massive log files when you update systems.

Because `/var` must be in place before a fresh installation of Red Hat Enterprise Linux CoreOS (RHCOS), the following procedure sets up the separate `/var` partition by creating a machine config manifest that is inserted during the **openshift-install** preparation phases of an OpenShift Container Platform installation.

Procedure

1. Create a directory to hold the OpenShift Container Platform installation files:

```
$ mkdir $HOME/clusterconfig
```

2. Run **openshift-install** to create a set of files in the **manifest** and **openshift** subdirectories. Answer the system questions as you are prompted:

```
$ openshift-install create manifests --dir $HOME/clusterconfig
? SSH Public Key ...
$ ls $HOME/clusterconfig/openshift/
99_kubeadmin-password-secret.yaml
99_openshift-cluster-api_master-machines-0.yaml
99_openshift-cluster-api_master-machines-1.yaml
99_openshift-cluster-api_master-machines-2.yaml
...
```

3. Create a Butane config that configures the additional partition. For example, name the file **`$HOME/clusterconfig/98-var-partition.bu`**, change the disk device name to the name of the storage device on the **worker** systems, and set the storage size as appropriate. This example places the `/var` directory on a separate partition:

```

variant: openshift
version: 4.16.0
metadata:
  labels:
    machineconfiguration.openshift.io/role: worker
  name: 98-var-partition
storage:
  disks:
    - device: /dev/disk/by-id/<device_name> ❶
      partitions:
        - label: var
          start_mib: <partition_start_offset> ❷
          size_mib: <partition_size> ❸
          number: 5
      filesystems:
        - device: /dev/disk/by-partlabel/var
          path: /var
          format: xfs
          mount_options: [defaults, prjquota] ❹
          with_mount_unit: true

```

- ❶ The storage device name of the disk that you want to partition.
- ❷ When adding a data partition to the boot disk, a minimum value of 25000 mebibytes is recommended. The root file system is automatically resized to fill all available space up to the specified offset. If no value is specified, or if the specified value is smaller than the recommended minimum, the resulting root file system will be too small, and future reinstalls of RHCOS might overwrite the beginning of the data partition.
- ❸ The size of the data partition in mebibytes.
- ❹ The **prjquota** mount option must be enabled for filesystems used for container storage.



NOTE

When creating a separate **/var** partition, you cannot use different instance types for worker nodes, if the different instance types do not have the same device name.

4. Create a manifest from the Butane config and save it to the **clusterconfig/openshift** directory. For example, run the following command:

```
$ butane $HOME/clusterconfig/98-var-partition.bu -o $HOME/clusterconfig/openshift/98-var-partition.yaml
```

5. Run **openshift-install** again to create Ignition configs from a set of files in the **manifest** and **openshift** subdirectories:

```
$ openshift-install create ignition-configs --dir $HOME/clusterconfig
$ ls $HOME/clusterconfig/
auth bootstrap.ign master.ign metadata.json worker.ign
```

Now you can use the Ignition config files as input to the vSphere installation procedures to install Red Hat Enterprise Linux CoreOS (RHCOS) systems.

23.3.4.13. Waiting for the bootstrap process to complete

The OpenShift Container Platform bootstrap process begins after the cluster nodes first boot into the persistent RHCOS environment that has been installed to disk. The configuration information provided through the Ignition config files is used to initialize the bootstrap process and install OpenShift Container Platform on the machines. You must wait for the bootstrap process to complete.

Prerequisites

- You have created the Ignition config files for your cluster.
- You have configured suitable network, DNS and load balancing infrastructure.
- You have obtained the installation program and generated the Ignition config files for your cluster.
- You installed RHCOS on your cluster machines and provided the Ignition config files that the OpenShift Container Platform installation program generated.
- Your machines have direct internet access or have an HTTP or HTTPS proxy available.

Procedure

1. Monitor the bootstrap process:

```
$ ./openshift-install --dir <installation_directory> wait-for bootstrap-complete \ 1  
--log-level=info 2
```

1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Example output

```
INFO Waiting up to 30m0s for the Kubernetes API at https://api.test.example.com:6443...  
INFO API v1.29.4 up  
INFO Waiting up to 30m0s for bootstrapping to complete...  
INFO It is now safe to remove the bootstrap resources
```

The command succeeds when the Kubernetes API server signals that it has been bootstrapped on the control plane machines.

2. After the bootstrap process is complete, remove the bootstrap machine from the load balancer.



IMPORTANT

You must remove the bootstrap machine from the load balancer at this point. You can also remove or reformat the bootstrap machine itself.

23.3.4.14. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

23.3.4.15. Approving the certificate signing requests for your machines

When you add machines to a cluster, two pending certificate signing requests (CSRs) are generated for each machine that you added. You must confirm that these CSRs are approved or, if necessary, approve them yourself. The client requests must be approved first, followed by the server requests.

Prerequisites

- You added machines to your cluster.

Procedure

1. Confirm that the cluster recognizes the machines:

```
$ oc get nodes
```

Example output

```
NAME      STATUS  ROLES  AGE  VERSION
master-0  Ready   master 63m  v1.29.4
master-1  Ready   master 63m  v1.29.4
master-2  Ready   master 64m  v1.29.4
```

The output lists all of the machines that you created.



NOTE

The preceding output might not include the compute nodes, also known as worker nodes, until some CSRs are approved.

- Review the pending CSRs and ensure that you see the client requests with the **Pending** or **Approved** status for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

```
NAME          AGE    REQUESTOR                                     CONDITION
csr-8b2br     15m   system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper  Pending
csr-8vnps     15m   system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper  Pending
...
```


In this example, two machines are joining the cluster. You might see more approved CSRs in the list.

- If the CSRs were not approved, after all of the pending CSRs for the machines you added are in **Pending** status, approve the CSRs for your cluster machines:



NOTE

Because the CSRs rotate automatically, approve your CSRs within an hour of adding the machines to the cluster. If you do not approve them within an hour, the certificates will rotate, and more than two certificates will be present for each node. You must approve all of these certificates. After the client CSR is approved, the Kubelet creates a secondary CSR for the serving certificate, which requires manual approval. Then, subsequent serving certificate renewal requests are automatically approved by the **machine-approver** if the Kubelet requests a new certificate with identical parameters.



NOTE

For clusters running on platforms that are not machine API enabled, such as bare metal and other user-provisioned infrastructure, you must implement a method of automatically approving the kubelet serving certificate requests (CSRs). If a request is not approved, then the **oc exec**, **oc rsh**, and **oc logs** commands cannot succeed, because a serving certificate is required when the API server connects to the kubelet. Any operation that contacts the Kubelet endpoint requires this certificate approval to be in place. The method must watch for new CSRs, confirm that the CSR was submitted by the **node-bootstrapper** service account in the **system:node** or **system:admin** groups, and confirm the identity of the node.


- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

1 **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{\n"}\n{{end}}' | xargs --no-run-if-empty oc adm certificate approve
```



NOTE

Some Operators might not become available until some CSRs are approved.

- Now that your client requests are approved, you must review the server requests for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

```
NAME      AGE   REQUESTOR                                     CONDITION
csr-bfd72 5m26s system:node:ip-10-0-50-126.us-east-2.compute.internal
Pending
csr-c57lv 5m26s system:node:ip-10-0-95-157.us-east-2.compute.internal
Pending
...
```

- If the remaining CSRs are not approved, and are in the **Pending** status, approve the CSRs for your cluster machines:

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

1 **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{\n"}\n{{end}}' | xargs oc adm certificate approve
```

- After all client and server CSRs have been approved, the machines have the **Ready** status. Verify this by running the following command:

```
$ oc get nodes
```

Example output

```
NAME      STATUS   ROLES    AGE   VERSION
master-0 Ready    master   73m   v1.29.4
```

```

master-1 Ready   master 73m v1.29.4
master-2 Ready   master 74m v1.29.4
worker-0 Ready   worker 11m v1.29.4
worker-1 Ready   worker 11m v1.29.4

```



NOTE

It can take a few minutes after approval of the server CSRs for the machines to transition to the **Ready** status.

Additional information

- For more information on CSRs, see [Certificate Signing Requests](#).

23.3.4.15.1. Initial Operator configuration

After the control plane initializes, you must immediately configure some Operators so that they all become available.

Prerequisites

- Your control plane has initialized.

Procedure

1. Watch the cluster components come online:

```
$ watch -n5 oc get clusteroperators
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.16.0	True	False	False	19m
baremetal	4.16.0	True	False	False	37m
cloud-credential	4.16.0	True	False	False	40m
cluster-autoscaler	4.16.0	True	False	False	37m
config-operator	4.16.0	True	False	False	38m
console	4.16.0	True	False	False	26m
csi-snapshot-controller	4.16.0	True	False	False	37m
dns	4.16.0	True	False	False	37m
etcd	4.16.0	True	False	False	36m
image-registry	4.16.0	True	False	False	31m
ingress	4.16.0	True	False	False	30m
insights	4.16.0	True	False	False	31m
kube-apiserver	4.16.0	True	False	False	26m
kube-controller-manager	4.16.0	True	False	False	36m
kube-scheduler	4.16.0	True	False	False	36m
kube-storage-version-migrator	4.16.0	True	False	False	37m
machine-api	4.16.0	True	False	False	29m
machine-approver	4.16.0	True	False	False	37m
machine-config	4.16.0	True	False	False	36m
marketplace	4.16.0	True	False	False	37m
monitoring	4.16.0	True	False	False	29m

network	4.16.0	True	False	False	38m
node-tuning	4.16.0	True	False	False	37m
openshift-apiserver	4.16.0	True	False	False	32m
openshift-controller-manager	4.16.0	True	False	False	30m
openshift-samples	4.16.0	True	False	False	32m
operator-lifecycle-manager	4.16.0	True	False	False	37m
operator-lifecycle-manager-catalog	4.16.0	True	False	False	37m
operator-lifecycle-manager-packageserver	4.16.0	True	False	False	32m
service-ca	4.16.0	True	False	False	38m
storage	4.16.0	True	False	False	37m

2. Configure the Operators that are not available.

23.3.4.15.2. Image registry removed during installation

On platforms that do not provide shareable object storage, the OpenShift Image Registry Operator bootstraps itself as **Removed**. This allows **openshift-installer** to complete installations on these platform types.

After installation, you must edit the Image Registry Operator configuration to switch the **managementState** from **Removed** to **Managed**. When this has completed, you must configure storage.

23.3.4.15.3. Image registry storage configuration

The Image Registry Operator is not initially available for platforms that do not provide default storage. After installation, you must configure your registry to use storage so that the Registry Operator is made available.

Instructions are shown for configuring a persistent volume, which is required for production clusters. Where applicable, instructions are shown for configuring an empty directory as the storage location, which is available for only non-production clusters.

Additional instructions are provided for allowing the image registry to use block storage types by using the **Recreate** rollout strategy during upgrades.

23.3.4.15.3.1. Configuring block registry storage for VMware vSphere

To allow the image registry to use block storage types such as vSphere Virtual Machine Disk (VMDK) during upgrades as a cluster administrator, you can use the **Recreate** rollout strategy.



IMPORTANT

Block storage volumes are supported but not recommended for use with image registry on production clusters. An installation where the registry is configured on block storage is not highly available because the registry cannot have more than one replica.

Procedure

1. Enter the following command to set the image registry storage as a block storage type, patch the registry so that it uses the **Recreate** rollout strategy, and runs with only **1** replica:

```
$ oc patch config.imageregistry.operator.openshift.io/cluster --type=merge -p '{"spec": {"rolloutStrategy": "Recreate", "replicas": 1}}'
```

2. Provision the PV for the block storage device, and create a PVC for that volume. The requested block volume uses the ReadWriteOnce (RWO) access mode.
 - a. Create a **pvc.yaml** file with the following contents to define a VMware vSphere **PersistentVolumeClaim** object:

```
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: image-registry-storage ❶
  namespace: openshift-image-registry ❷
spec:
  accessModes:
  - ReadWriteOnce ❸
  resources:
    requests:
      storage: 100Gi ❹
```

- ❶ A unique name that represents the **PersistentVolumeClaim** object.
- ❷ The namespace for the **PersistentVolumeClaim** object, which is **openshift-image-registry**.
- ❸ The access mode of the persistent volume claim. With **ReadWriteOnce**, the volume can be mounted with read and write permissions by a single node.
- ❹ The size of the persistent volume claim.

- b. Enter the following command to create the **PersistentVolumeClaim** object from the file:

```
$ oc create -f pvc.yaml -n openshift-image-registry
```

3. Enter the following command to edit the registry configuration so that it references the correct PVC:

```
$ oc edit config.imageregistry.operator.openshift.io -o yaml
```

Example output

```
storage:
  pvc:
    claim: ❶
```

- ❶ By creating a custom PVC, you can leave the **claim** field blank for the default automatic creation of an **image-registry-storage** PVC.

For instructions about configuring registry storage so that it references the correct PVC, see [Configuring the registry for vSphere](#).

23.3.4.16. Completing installation on user-provisioned infrastructure

After you complete the Operator configuration, you can finish installing the cluster on infrastructure that you provide.

Prerequisites

- Your control plane has initialized.
- You have completed the initial Operator configuration.

Procedure

1. Confirm that all the cluster components are online with the following command:

```
$ watch -n5 oc get clusteroperators
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.16.0	True	False	False	19m
baremetal	4.16.0	True	False	False	37m
cloud-credential	4.16.0	True	False	False	40m
cluster-autoscaler	4.16.0	True	False	False	37m
config-operator	4.16.0	True	False	False	38m
console	4.16.0	True	False	False	26m
csi-snapshot-controller	4.16.0	True	False	False	37m
dns	4.16.0	True	False	False	37m
etcd	4.16.0	True	False	False	36m
image-registry	4.16.0	True	False	False	31m
ingress	4.16.0	True	False	False	30m
insights	4.16.0	True	False	False	31m
kube-apiserver	4.16.0	True	False	False	26m
kube-controller-manager	4.16.0	True	False	False	36m
kube-scheduler	4.16.0	True	False	False	36m
kube-storage-version-migrator	4.16.0	True	False	False	37m
machine-api	4.16.0	True	False	False	29m
machine-approver	4.16.0	True	False	False	37m
machine-config	4.16.0	True	False	False	36m
marketplace	4.16.0	True	False	False	37m
monitoring	4.16.0	True	False	False	29m
network	4.16.0	True	False	False	38m
node-tuning	4.16.0	True	False	False	37m
openshift-apiserver	4.16.0	True	False	False	32m
openshift-controller-manager	4.16.0	True	False	False	30m
openshift-samples	4.16.0	True	False	False	32m
operator-lifecycle-manager	4.16.0	True	False	False	37m
operator-lifecycle-manager-catalog	4.16.0	True	False	False	37m
operator-lifecycle-manager-packageserver	4.16.0	True	False	False	32m
service-ca	4.16.0	True	False	False	38m
storage	4.16.0	True	False	False	37m

Alternatively, the following command notifies you when all of the clusters are available. It also retrieves and displays credentials:

```
$ ./openshift-install --dir <installation_directory> wait-for install-complete 1
```

-
- 1 For `<installation_directory>`, specify the path to the directory that you stored the installation files in.

Example output

```
INFO Waiting up to 30m0s for the cluster to initialize...
```

The command succeeds when the Cluster Version Operator finishes deploying the OpenShift Container Platform cluster from Kubernetes API server.



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

2. Confirm that the Kubernetes API server is communicating with the pods.
 - a. To view a list of all pods, use the following command:

```
$ oc get pods --all-namespaces
```

Example output

```

NAMESPACE           NAME                                     READY  STATUS
RESTARTS  AGE
openshift-apiserver-operator  openshift-apiserver-operator-85cb746d55-zqhs8  1/1
Running   1    9m
openshift-apiserver          apiserver-67b9g                                1/1  Running  0
3m
openshift-apiserver          apiserver-ljcmx                                1/1  Running  0
1m
openshift-apiserver          apiserver-z25h4                                1/1  Running  0
2m
openshift-authentication-operator  authentication-operator-69d5d8bf84-vh2n8  1/1
Running   0    5m
...

```

- b. View the logs for a pod that is listed in the output of the previous command by using the following command:

```
$ oc logs <pod_name> -n <namespace> 1
```

- 1 Specify the pod name and namespace, as shown in the output of the previous command.

If the pod logs display, the Kubernetes API server can communicate with the cluster machines.

3. For an installation with Fibre Channel Protocol (FCP), additional steps are required to enable multipathing. Do not enable multipathing during installation. See "Enabling multipathing with kernel arguments on RHCOS" in the *Postinstallation machine configuration tasks* documentation for more information.

You can add extra compute machines after the cluster installation is completed by following [Adding compute machines to vSphere](#).

23.3.4.17. Configuring vSphere DRS anti-affinity rules for control plane nodes

vSphere Distributed Resource Scheduler (DRS) anti-affinity rules can be configured to support higher availability of OpenShift Container Platform Control Plane nodes. Anti-affinity rules ensure that the vSphere Virtual Machines for the OpenShift Container Platform Control Plane nodes are not scheduled to the same vSphere Host.



IMPORTANT

- The following information applies to compute DRS only and does not apply to storage DRS.
- The **govc** command is an open-source command available from VMware; it is not available from Red Hat. The **govc** command is not supported by the Red Hat support.
- Instructions for downloading and installing **govc** are found on the VMware documentation website.

Create an anti-affinity rule by running the following command:

Example command

```
$ govc cluster.rule.create \
  -name openshift4-control-plane-group \
  -dc MyDatacenter -cluster MyCluster \
  -enable \
  -anti-affinity master-0 master-1 master-2
```

After creating the rule, your control plane nodes are automatically migrated by vSphere so they are not running on the same hosts. This might take some time while vSphere reconciles the new rule. Successful command completion is shown in the following procedure.



NOTE

The migration occurs automatically and might cause brief OpenShift API outage or latency until the migration finishes.

The vSphere DRS anti-affinity rules need to be updated manually in the event of a control plane VM name change or migration to a new vSphere Cluster.

Procedure

1. Remove any existing DRS anti-affinity rule by running the following command:

```
$ govc cluster.rule.remove \  
-name openshift4-control-plane-group \  
-dc MyDatacenter -cluster MyCluster
```

Example Output

```
[13-10-22 09:33:24] Reconfigure /MyDatacenter/host/MyCluster...OK
```

2. Create the rule again with updated names by running the following command:

```
$ govc cluster.rule.create \  
-name openshift4-control-plane-group \  
-dc MyDatacenter -cluster MyOtherCluster \  
-enable \  
-anti-affinity master-0 master-1 master-2
```

23.3.4.18. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service

23.3.4.19. Next steps

- [Customize your cluster](#).
- If necessary, you can [opt out of remote health reporting](#).
- [Set up your registry and configure registry storage](#).
- Optional: [View the events from the vSphere Problem Detector Operator](#) to determine if the cluster has permission or storage configuration issues.
- Optional: if you created encrypted virtual machines, [create an encrypted storage class](#).

23.3.5. Installing a cluster on vSphere in a restricted network with user-provisioned infrastructure

In OpenShift Container Platform version 4.16, you can install a cluster on VMware vSphere infrastructure that you provision in a restricted network.



NOTE

OpenShift Container Platform supports deploying a cluster to a single VMware vCenter only. Deploying a cluster with machines/machine sets on multiple vCenters is not supported.




IMPORTANT

The steps for performing a user-provisioned infrastructure installation are provided as an example only. Installing a cluster with infrastructure you provide requires knowledge of the vSphere platform and the installation process of OpenShift Container Platform. Use the user-provisioned infrastructure installation instructions as a guide; you are free to create the required resources through other methods.

23.3.5.1. Prerequisites


- You have completed the tasks in [Preparing to install a cluster using user-provisioned infrastructure](#).
- You reviewed your VMware platform licenses. Red Hat does not place any restrictions on your VMware licenses, but some VMware infrastructure components require licensing.
- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- You [created a registry on your mirror host](#) and obtained the **imageContentSources** data for your version of OpenShift Container Platform.



IMPORTANT

Because the installation media is on the mirror host, you can use that computer to complete all installation steps.

- You provisioned [persistent storage](#) for your cluster. To deploy a private image registry, your storage must provide **ReadWriteMany** access modes.
- Completing the installation requires that you upload the Red Hat Enterprise Linux CoreOS (RHCOS) OVA on vSphere hosts. The machine from which you complete this process requires access to port 443 on the vCenter and ESXi hosts. You verified that port 443 is accessible.
- If you use a firewall, you confirmed with the administrator that port 443 is accessible. Control plane nodes must be able to reach vCenter and ESXi hosts on port 443 for the installation to succeed.
- If you use a firewall and plan to use the Telemetry service, you [configured the firewall to allow the sites](#) that your cluster requires access to.

**NOTE**


Be sure to also review this site list if you are configuring a proxy.

23.3.5.2. About installations in restricted networks

In OpenShift Container Platform 4.16, you can perform an installation that does not require an active connection to the internet to obtain software components. Restricted network installations can be completed using installer-provisioned infrastructure or user-provisioned infrastructure, depending on the cloud platform to which you are installing the cluster.

If you choose to perform a restricted network installation on a cloud platform, you still require access to its cloud APIs. Some cloud functions, like Amazon Web Service's Route 53 DNS and IAM services, require internet access. Depending on your network, you might require less internet access for an installation on bare metal hardware, Nutanix, or on VMware vSphere.

To complete a restricted network installation, you must create a registry that mirrors the contents of the OpenShift image registry and contains the installation media. You can create this registry on a mirror host, which can access both the internet and your closed network, or by using other methods that meet your restrictions.

**IMPORTANT**

Because of the complexity of the configuration for user-provisioned installations, consider completing a standard user-provisioned infrastructure installation before you attempt a restricted network installation using user-provisioned infrastructure. Completing this test installation might make it easier to isolate and troubleshoot any issues that might arise during your installation in a restricted network.

23.3.5.2.1. Additional limits

Clusters in restricted networks have the following additional limitations and restrictions:

- The **ClusterVersion** status includes an **Unable to retrieve available updates** error.
- By default, you cannot use the contents of the Developer Catalog because you cannot access the required image stream tags.

23.3.5.3. Internet access for OpenShift Container Platform


In OpenShift Container Platform 4.16, you require access to the internet to obtain the images that are necessary to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.

23.3.5.4. VMware vSphere region and zone enablement

You can deploy an OpenShift Container Platform cluster to multiple vSphere datacenters that run in a single VMware vCenter. Each datacenter can run multiple clusters. This configuration reduces the risk of a hardware failure or network outage that can cause your cluster to fail. To enable regions and zones, you must define multiple failure domains for your OpenShift Container Platform cluster.



IMPORTANT

The VMware vSphere region and zone enablement feature requires the vSphere Container Storage Interface (CSI) driver as the default storage driver in the cluster. As a result, the feature is only available on a newly installed cluster.

For a cluster that was upgraded from a previous release, you must enable CSI automatic migration for the cluster. You can then configure multiple regions and zones for the upgraded cluster.

The default installation configuration deploys a cluster to a single vSphere datacenter. If you want to deploy a cluster to multiple vSphere datacenters, you must create an installation configuration file that enables the region and zone feature.

The default **install-config.yaml** file includes **vccenters** and **failureDomains** fields, where you can specify multiple vSphere datacenters and clusters for your OpenShift Container Platform cluster. You can leave these fields blank if you want to install an OpenShift Container Platform cluster in a vSphere environment that consists of single datacenter.

The following list describes terms associated with defining zones and regions for your cluster:

- **Failure domain:** Establishes the relationships between a region and zone. You define a failure domain by using vCenter objects, such as a **datastore** object. A failure domain defines the vCenter location for OpenShift Container Platform cluster nodes.
- **Region:** Specifies a vCenter datacenter. You define a region by using a tag from the **openshift-region** tag category.
- **Zone:** Specifies a vCenter cluster. You define a zone by using a tag from the **openshift-zone** tag category.



NOTE

If you plan on specifying more than one failure domain in your **install-config.yaml** file, you must create tag categories, zone tags, and region tags in advance of creating the configuration file.

You must create a vCenter tag for each vCenter datacenter, which represents a region. Additionally, you must create a vCenter tag for each cluster that runs in a datacenter, which represents a zone. After you create the tags, you must attach each tag to their respective datacenters and clusters.

The following table outlines an example of the relationship among regions, zones, and tags for a configuration with multiple vSphere datacenters running in a single VMware vCenter.

Datacenter (region)	Cluster (zone)	Tags
us-east	us-east-1	us-east-1a


Datacenter (region)	Cluster (zone)	Tags
	us-east-2	us-east-1b
		us-east-2a
		us-east-2b
us-west	us-west-1	us-west-1a
		us-west-1b
	us-west-2	us-west-2a
		us-west-2b

Additional resources

- [Additional VMware vSphere configuration parameters](#)
- [Deprecated VMware vSphere configuration parameters](#)
- [vSphere automatic migration](#)
- [VMware vSphere CSI Driver Operator](#)

23.3.5.5. Manually creating the installation configuration file

Installing the cluster requires that you manually create the installation configuration file.



IMPORTANT

The Cloud Controller Manager Operator performs a connectivity check on a provided hostname or IP address. Ensure that you specify a hostname or an IP address to a reachable vCenter server. If you provide metadata to a non-existent vCenter server, installation of the cluster fails at the bootstrap stage.


Prerequisites

- You have an SSH public key on your local machine to provide to the installation program. The key will be used for SSH authentication onto your cluster nodes for debugging and disaster recovery.
- You have obtained the OpenShift Container Platform installation program and the pull secret for your cluster.
- Obtain the **imageContentSources** section from the output of the command to mirror the repository.
- Obtain the contents of the certificate for your mirror registry.

Procedure

1. Create an installation directory to store your required installation assets in:

```
$ mkdir <installation_directory>
```



IMPORTANT

You must create a directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.


2. Customize the sample **install-config.yaml** file template that is provided and save it in the **<installation_directory>**.



NOTE

You must name this configuration file **install-config.yaml**.

- Unless you use a registry that RHCOS trusts by default, such as **docker.io**, you must provide the contents of the certificate for your mirror repository in the **additionalTrustBundle** section. In most cases, you must provide the certificate for your mirror.
- You must include the **imageContentSources** section from the output of the command to mirror the repository.



IMPORTANT

- The **ImageContentSourcePolicy** file is generated as an output of **oc mirror** after the mirroring process is finished.
- The **oc mirror** command generates an **ImageContentSourcePolicy** file which contains the YAML needed to define **ImageContentSourcePolicy**. Copy the text from this file and paste it into your **install-config.yaml** file.
- You must run the 'oc mirror' command twice. The first time you run the **oc mirror** command, you get a full **ImageContentSourcePolicy** file. The second time you run the **oc mirror** command, you only get the difference between the first and second run. Because of this behavior, you must always keep a backup of these files in case you need to merge them into one complete **ImageContentSourcePolicy** file. Keeping a backup of these two output files ensures that you have a complete **ImageContentSourcePolicy** file.

3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



IMPORTANT

The **install-config.yaml** file is consumed during the next step of the installation process. You must back it up now.

Additional resources

- [Installation configuration parameters](#)

23.3.5.5.1. Sample `install-config.yaml` file for VMware vSphere

You can customize the `install-config.yaml` file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.

```

additionalTrustBundlePolicy: Proxyonly
apiVersion: v1
baseDomain: example.com 1
compute: 2
- architecture: amd64
  name: <worker_node>
  platform: {}
  replicas: 0 3
controlPlane: 4
  architecture: amd64
  name: <parent_node>
  platform: {}
  replicas: 3 5
metadata:
  creationTimestamp: null
  name: test 6
networking:
---
platform:
  vsphere:
    failureDomains: 7
    - name: <failure_domain_name>
      region: <default_region_name>
      server: <fully_qualified_domain_name>
    topology:
      computeCluster: "/<datacenter>/host/<cluster>"
      datacenter: <datacenter> 8
      datastore: "/<datacenter>/datastore/<datastore>" 9
      networks:
      - <VM_Network_name>
      resourcePool: "/<datacenter>/host/<cluster>/Resources/<resourcePool>" 10
      folder: "/<datacenter_name>/vm/<folder_name>/<subfolder_name>" 11
      zone: <default_zone_name>
    vcenters:
    - datacenters:
      - <datacenter>
      password: <password> 12
      port: 443
      server: <fully_qualified_domain_name> 13
      user: administrator@vsphere.local
      diskType: thin 14
    fips: false 15
  pullSecret: '{"auths":{"<local_registry>":{"auth": "<credentials>","email": "you@example.com"}}}' 16
  sshKey: 'ssh-ed25519 AAAA..' 17

```


- 11 Optional: For installer-provisioned infrastructure, the absolute path of an existing folder where the installation program creates the virtual machines, for example, `/<datacenter_name>/vm/<folder_name>/<subfolder_name>`. If you do not provide this value, the installation program creates a top-level folder in the datacenter virtual machine folder that is named with the infrastructure ID. If you are providing the infrastructure for the cluster and you do not want to use the default **StorageClass** object, named **thin**, you can omit the **folder** parameter from the **install-config.yaml** file.
- 12 The password associated with the vSphere user.
- 13 The fully-qualified hostname or IP address of the vCenter server.



IMPORTANT

The Cloud Controller Manager Operator performs a connectivity check on a provided hostname or IP address. Ensure that you specify a hostname or an IP address to a reachable vCenter server. If you provide metadata to a non-existent vCenter server, installation of the cluster fails at the bootstrap stage.

- 14 The vSphere disk provisioning method.
- 15 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.



IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 16 For `<local_registry>`, specify the registry domain name, and optionally the port, that your mirror registry uses to serve content. For example **registry.example.com** or **registry.example.com:5000**. For `<credentials>`, specify the base64-encoded user name and password for your mirror registry.
- 17 The public portion of the default SSH key for the **core** user in Red Hat Enterprise Linux CoreOS (RHCOS).



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- 18 Provide the contents of the certificate file that you used for your mirror registry.

- 19 Provide the **imageContentSources** section from the output of the command to mirror the repository.

23.3.5.5.2. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

- Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- A proxy URL to use for creating HTTPS connections outside the cluster.
- A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all

destinations. You must include vCenter’s IP address and the IP range that you use for its machines.

- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy’s identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.



NOTE

The installation program does not support the proxy **readinessEndpoints** field.



NOTE

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$ ./openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.



NOTE

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

23.3.5.5.3. Configuring regions and zones for a VMware vCenter

You can modify the default installation configuration file, so that you can deploy an OpenShift Container Platform cluster to multiple vSphere datacenters that run in a single VMware vCenter.

The default **install-config.yaml** file configuration from the previous release of OpenShift Container Platform is deprecated. You can continue to use the deprecated default configuration, but the **openshift-installer** will prompt you with a warning message that indicates the use of deprecated fields in the configuration file.



IMPORTANT

The example uses the **govc** command. The **govc** command is an open source command available from VMware; it is not available from Red Hat. The Red Hat support team does not maintain the **govc** command. Instructions for downloading and installing **govc** are found on the VMware documentation website

Prerequisites

- You have an existing **install-config.yaml** installation configuration file.




IMPORTANT

You must specify at least one failure domain for your OpenShift Container Platform cluster, so that you can provision datacenter objects for your VMware vCenter server. Consider specifying multiple failure domains if you need to provision virtual machine nodes in different datacenters, clusters, datastores, and other components. To enable regions and zones, you must define multiple failure domains for your OpenShift Container Platform cluster.

Procedure

1. Enter the following **govc** command-line tool commands to create the **openshift-region** and **openshift-zone** vCenter tag categories:



IMPORTANT

If you specify different names for the **openshift-region** and **openshift-zone** vCenter tag categories, the installation of the OpenShift Container Platform cluster fails.

```
$ govc tags.category.create -d "OpenShift region" openshift-region
```

```
$ govc tags.category.create -d "OpenShift zone" openshift-zone
```

2. To create a region tag for each region vSphere datacenter where you want to deploy your cluster, enter the following command in your terminal:

```
$ govc tags.create -c <region_tag_category> <region_tag>
```

3. To create a zone tag for each vSphere cluster where you want to deploy your cluster, enter the following command:

```
$ govc tags.create -c <zone_tag_category> <zone_tag>
```

4. Attach region tags to each vCenter datacenter object by entering the following command:

```
$ govc tags.attach -c <region_tag_category> <region_tag_1> /<datacenter_1>
```

5. Attach the zone tags to each vCenter datacenter object by entering the following command:

```
$ govc tags.attach -c <zone_tag_category> <zone_tag_1> /<datacenter_1>/host/vcs-mdcnc-workload-1
```

6. Change to the directory that contains the installation program and initialize the cluster deployment according to your chosen installation requirements.

Sample install-config.yaml file with multiple datacenters defined in a vSphere center

```
---
compute:
---
vsphere:
  zones:
    - "<machine_pool_zone_1>"
    - "<machine_pool_zone_2>"
---
controlPlane:
---
vsphere:
  zones:
    - "<machine_pool_zone_1>"
    - "<machine_pool_zone_2>"
---
platform:
  vsphere:
    vcenters:
---
  datacenters:
    - <datacenter1_name>
    - <datacenter2_name>
  failureDomains:
    - name: <machine_pool_zone_1>
      region: <region_tag_1>
      zone: <zone_tag_1>
      server: <fully_qualified_domain_name>
      topology:
        datacenter: <datacenter1>
        computeCluster: "/<datacenter1>/host/<cluster1>"
        networks:
          - <VM_Network1_name>
        datastore: "/<datacenter1>/datastore/<datastore1>"
        resourcePool: "/<datacenter1>/host/<cluster1>/Resources/<resourcePool1>"
        folder: "/<datacenter1>/vm/<folder1>"
    - name: <machine_pool_zone_2>
      region: <region_tag_2>
      zone: <zone_tag_2>
      server: <fully_qualified_domain_name>
      topology:
        datacenter: <datacenter2>
        computeCluster: "/<datacenter2>/host/<cluster2>"
        networks:
          - <VM_Network2_name>
        datastore: "/<datacenter2>/datastore/<datastore2>"
```


```
resourcePool: "/<datacenter2>/host/<cluster2>/Resources/<resourcePool2>"
folder: "/<datacenter2>/vm/<folder2>"
```

```
---
```

23.3.5.6. Creating the Kubernetes manifest and Ignition config files

Because you must modify some cluster definition files and manually start the cluster machines, you must generate the Kubernetes manifest and Ignition config files that the cluster needs to configure the machines.

The installation configuration file transforms into the Kubernetes manifests. The manifests wrap into the Ignition configuration files, which are later used to configure the cluster machines.



IMPORTANT

- The Ignition config files that the OpenShift Container Platform installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

Prerequisites

- You obtained the OpenShift Container Platform installation program. For a restricted network installation, these files are on your mirror host.
- You created the **install-config.yaml** installation configuration file.

Procedure

1. Change to the directory that contains the OpenShift Container Platform installation program and generate the Kubernetes manifests for the cluster:

```
$ ./openshift-install create manifests --dir <installation_directory> 1
```

- 1** For **<installation_directory>**, specify the installation directory that contains the **install-config.yaml** file you created.

2. Remove the Kubernetes manifest files that define the control plane machines, compute machine sets, and control plane machine sets:

```
$ rm -f openshift/99_openshift-cluster-api_master-machines-*.yaml openshift/99_openshift-cluster-api_worker-machineset-*.yaml openshift/99_openshift-machine-api_master-control-plane-machine-set.yaml
```

Because you create and manage these resources yourself, you do not have to initialize them.

- You can preserve the compute machine set files to create compute machines by using the machine API, but you must update references to them to match your environment.
3. Check that the **mastersSchedulable** parameter in the `<installation_directory>/manifests/cluster-scheduler-02-config.yml` Kubernetes manifest file is set to **false**. This setting prevents pods from being scheduled on the control plane machines:
 - a. Open the `<installation_directory>/manifests/cluster-scheduler-02-config.yml` file.
 - b. Locate the **mastersSchedulable** parameter and ensure that it is set to **false**.
 - c. Save and exit the file.
 4. To create the Ignition configuration files, run the following command from the directory that contains the installation program:

```
$ ./openshift-install create ignition-configs --dir <installation_directory> 1
```

- 1** For `<installation_directory>`, specify the same installation directory.

Ignition config files are created for the bootstrap, control plane, and compute nodes in the installation directory. The **kubeadmin-password** and **kubeconfig** files are created in the `./<installation_directory>/auth` directory:


```
.
├── auth
│   ├── kubeadmin-password
│   └── kubeconfig
├── bootstrap.ign
├── master.ign
├── metadata.json
└── worker.ign
```

23.3.5.7. Configuring chrony time service

You must set the time server and related settings used by the chrony time service (**chronyd**) by modifying the contents of the **chrony.conf** file and passing those contents to your nodes as a machine config.

Procedure

1. Create a Butane config including the contents of the **chrony.conf** file. For example, to configure chrony on worker nodes, create a **99-worker-chrony.bu** file.



NOTE

See "Creating machine configs with Butane" for information about Butane.

```
variant: openshift
version: 4.16.0
```

```

metadata:
  name: 99-worker-chrony ❶
  labels:
    machineconfiguration.openshift.io/role: worker ❷
storage:
  files:
  - path: /etc/chrony.conf
    mode: 0644 ❸
    overwrite: true
    contents:
      inline: |
        pool 0.rhel.pool.ntp.org iburst ❹
        driftfile /var/lib/chrony/drift
        makestep 1.0 3
        rtcsync
        logdir /var/log/chrony

```

❶ ❷ On control plane nodes, substitute **master** for **worker** in both of these locations.

❸ Specify an octal value mode for the **mode** field in the machine config file. After creating the file and applying the changes, the **mode** is converted to a decimal value. You can check the YAML file with the command **oc get mc <mc-name> -o yaml**.

❹ Specify any valid, reachable time source, such as the one provided by your DHCP server.

- Use Butane to generate a **MachineConfig** object file, **99-worker-chrony.yaml**, containing the configuration to be delivered to the nodes:

```
$ butane 99-worker-chrony.bu -o 99-worker-chrony.yaml
```

- Apply the configurations in one of two ways:

- If the cluster is not running yet, after you generate manifest files, add the **MachineConfig** object file to the **<installation_directory>/openshift** directory, and then continue to create the cluster.
- If the cluster is already running, apply the file:

```
$ oc apply -f ./99-worker-chrony.yaml
```

23.3.5.8. Extracting the infrastructure name

The Ignition config files contain a unique cluster identifier that you can use to uniquely identify your cluster in VMware vSphere. If you plan to use the cluster identifier as the name of your virtual machine folder, you must extract it.

Prerequisites

- You obtained the OpenShift Container Platform installation program and the pull secret for your cluster.
- You generated the Ignition config files for your cluster.
- You installed the **jq** package.

Procedure

- To extract and view the infrastructure name from the Ignition config file metadata, run the following command:

```
$ jq -r .infraID <installation_directory>/metadata.json 1
```

- For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

Example output

```
openshift-vw9j6 1
```

- The output of this command is your cluster name and a random string.

23.3.5.9. Installing RHCOS and starting the OpenShift Container Platform bootstrap process

To install OpenShift Container Platform on user-provisioned infrastructure on VMware vSphere, you must install Red Hat Enterprise Linux CoreOS (RHCOS) on vSphere hosts. When you install RHCOS, you must provide the Ignition config file that was generated by the OpenShift Container Platform installation program for the type of machine you are installing. If you have configured suitable networking, DNS, and load balancing infrastructure, the OpenShift Container Platform bootstrap process begins automatically after the RHCOS machines have rebooted.

Prerequisites

- You have obtained the Ignition config files for your cluster.
- You have access to an HTTP server that you can access from your computer and that the machines that you create can access.
- You have created a [vSphere cluster](#).

Procedure

- Upload the bootstrap Ignition config file, which is named **<installation_directory>/bootstrap.ign**, that the installation program created to your HTTP server. Note the URL of this file.
- Save the following secondary Ignition config file for your bootstrap node to your computer as **<installation_directory>/merge-bootstrap.ign**:

```
{
  "ignition": {
    "config": {
      "merge": [
        {
          "source": "<bootstrap_ignition_config_url>", 1
        }
      ]
    }
  }
}
```

```

    ]
  },
  "timeouts": {},
  "version": "3.2.0"
},
"networkd": {},
"passwd": {},
"storage": {},
"systemd": {}
}

```

- 1 Specify the URL of the bootstrap Ignition config file that you hosted.


When you create the virtual machine (VM) for the bootstrap machine, you use this Ignition config file.

3. Locate the following Ignition config files that the installation program created:
 - `<installation_directory>/master.ign`
 - `<installation_directory>/worker.ign`
 - `<installation_directory>/merge-bootstrap.ign`
4. Convert the Ignition config files to Base64 encoding. Later in this procedure, you must add these files to the extra configuration parameter **guestinfo.ignition.config.data** in your VM. For example, if you use a Linux operating system, you can use the **base64** command to encode the files.

```
$ base64 -w0 <installation_directory>/master.ign > <installation_directory>/master.64
```

```
$ base64 -w0 <installation_directory>/worker.ign > <installation_directory>/worker.64
```


```
$ base64 -w0 <installation_directory>/merge-bootstrap.ign > <installation_directory>/merge-bootstrap.64
```



IMPORTANT

If you plan to add more compute machines to your cluster after you finish installation, do not delete these files.

5. Obtain the RHCOS OVA image. Images are available from the [RHCOS image mirror](#) page.



IMPORTANT

The RHCOS images might not change with every release of OpenShift Container Platform. You must download an image with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Use the image version that matches your OpenShift Container Platform version if it is available.

The filename contains the OpenShift Container Platform version number in the format **rhcos-vmware.<architecture>.ova**.

6. In the vSphere Client, create a folder in your datacenter to store your VMs.
 - a. Click the **VMs and Templates** view.
 - b. Right-click the name of your datacenter.
 - c. Click **New Folder → New VM and Template Folder**.
 - d. In the window that is displayed, enter the folder name. If you did not specify an existing folder in the **install-config.yaml** file, then create a folder with the same name as the infrastructure ID. You use this folder name so vCenter dynamically provisions storage in the appropriate location for its Workspace configuration.
7. In the vSphere Client, create a template for the OVA image and then clone the template as needed.



NOTE

In the following steps, you create a template and then clone the template for all of your cluster machines. You then provide the location for the Ignition config file for that cloned machine type when you provision the VMs.


- a. From the **Hosts and Clusters** tab, right-click your cluster name and select **Deploy OVF Template**.
- b. On the **Select an OVF** tab, specify the name of the RHCOS OVA file that you downloaded.
- c. On the **Select a name and folder** tab, set a **Virtual machine name** for your template, such as **Template-RHCOS**. Click the name of your vSphere cluster and select the folder you created in the previous step.
- d. On the **Select a compute resource** tab, click the name of your vSphere cluster.
- e. On the **Select storage** tab, configure the storage options for your VM.
 - Select **Thin Provision** or **Thick Provision**, based on your storage preferences.
 - Select the datastore that you specified in your **install-config.yaml** file.
 - If you want to encrypt your virtual machines, select **Encrypt this virtual machine**. See the section titled "Requirements for encrypting virtual machines" for more information.
- f. On the **Select network** tab, specify the network that you configured for the cluster, if available.
- g. When creating the OVF template, do not specify values on the **Customize template** tab or configure the template any further.



IMPORTANT

Do not start the original VM template. The VM template must remain off and must be cloned for new RHCOS machines. Starting the VM template configures the VM template as a VM on the platform, which prevents it from being used as a template that compute machine sets can apply configurations to.

8. Optional: Update the configured virtual hardware version in the VM template, if necessary. Follow [Upgrading a virtual machine to the latest hardware version](#) in the VMware documentation for more information.



IMPORTANT

It is recommended that you update the hardware version of the VM template to version 15 before creating VMs from it, if necessary. Using hardware version 13 for your cluster nodes running on vSphere is now deprecated. If your imported template defaults to hardware version 13, you must ensure that your ESXi host is on 6.7U3 or later before upgrading the VM template to hardware version 15. If your vSphere version is less than 6.7U3, you can skip this upgrade step; however, a future version of OpenShift Container Platform is scheduled to remove support for hardware version 13 and vSphere versions less than 6.7U3.


9. After the template deploys, deploy a VM for a machine in the cluster.
 - a. Right-click the template name and click **Clone → Clone to Virtual Machine**
 - b. On the **Select a name and folder** tab, specify a name for the VM. You might include the machine type in the name, such as **control-plane-0** or **compute-1**.



NOTE

Ensure that all virtual machine names across a vSphere installation are unique.

- c. On the **Select a name and folder** tab, select the name of the folder that you created for the cluster.
- d. On the **Select a compute resource** tab, select the name of a host in your datacenter.
- e. On the **Select clone options** tab, select **Customize this virtual machine's hardware**
- f. On the **Customize hardware** tab, click **Advanced Parameters**.



IMPORTANT

The following configuration suggestions are for example purposes only. As a cluster administrator, you must configure resources according to the resource demands placed on your cluster. To best manage cluster resources, consider creating a resource pool from the cluster's root resource pool.

- Optional: Override default DHCP networking in vSphere. To enable static IP networking:
 - Set your static IP configuration:

Example command

```
$ export IPCFG="ip=<ip>::<gateway>:<netmask>:<hostname>:<iface>:none
nameserver=svr1 [nameserver=svr2 [nameserver=svr3 [...]]]"
```

Example command

```
$ export IPCFG="ip=192.168.100.101::192.168.100.254:255.255.255.0::none
nameserver=8.8.8.8"
```

- Set the **guestinfo.afterburn.initrd.network-kargs** property before you boot a VM from an OVA in vSphere:

Example command

```
$ govc vm.change -vm "<vm_name>" -e "guestinfo.afterburn.initrd.network-
kargs=${IPCFG}"
```


- Add the following configuration parameter names and values by specifying data in the **Attribute** and **Values** fields. Ensure that you select the **Add** button for each parameter that you create.
 - **guestinfo.ignition.config.data**: Locate the base-64 encoded files that you created previously in this procedure, and paste the contents of the base64-encoded Ignition config file for this machine type.
 - **guestinfo.ignition.config.data.encoding**: Specify **base64**.
 - **disk.EnableUUID**: Specify **TRUE**.
 - **stealclock.enable**: If this parameter was not defined, add it and specify **TRUE**.
 - Create a child resource pool from the cluster's root resource pool. Perform resource allocation in this child resource pool.
- g. In the **Virtual Hardware** panel of the **Customize hardware** tab, modify the specified values as required. Ensure that the amount of RAM, CPU, and disk storage meets the minimum requirements for the machine type.
- h. Complete the remaining configuration steps. On clicking the **Finish** button, you have completed the cloning operation.
- i. From the **Virtual Machines** tab, right-click on your VM and then select **Power → Power On**.
- j. Check the console output to verify that Ignition ran.

Example command

```
Ignition: ran on 2022/03/14 14:48:33 UTC (this boot)
Ignition: user-provided config was applied
```

Next steps

- Create the rest of the machines for your cluster by following the preceding steps for each machine.



IMPORTANT

You must create the bootstrap and control plane machines at this time. Because some pods are deployed on compute machines by default, also create at least two compute machines before you install the cluster.

23.3.5.10. Adding more compute machines to a cluster in vSphere

You can add more compute machines to a user-provisioned OpenShift Container Platform cluster on VMware vSphere.

After your vSphere template deploys in your OpenShift Container Platform cluster, you can deploy a virtual machine (VM) for a machine in that cluster.

Prerequisites

- Obtain the base64-encoded Ignition file for your compute machines.
- You have access to the vSphere template that you created for your cluster.

Procedure

1. Right-click the template's name and click **Clone** → **Clone to Virtual Machine**
2. On the **Select a name and folder** tab, specify a name for the VM. You might include the machine type in the name, such as **compute-1**.



NOTE

Ensure that all virtual machine names across a vSphere installation are unique.

3. On the **Select a name and folder** tab, select the name of the folder that you created for the cluster.
4. On the **Select a compute resource** tab, select the name of a host in your datacenter.
5. On the **Select storage** tab, select storage for your configuration and disk files.
6. On the **Select clone options** tab, select **Customize this virtual machine's hardware**
7. On the **Customize hardware** tab, click **Advanced Parameters**.
 - Add the following configuration parameter names and values by specifying data in the **Attribute** and **Values** fields. Ensure that you select the **Add** button for each parameter that you create.
 - **guestinfo.ignition.config.data**: Paste the contents of the base64-encoded compute Ignition config file for this machine type.
 - **guestinfo.ignition.config.data.encoding**: Specify **base64**.
 - **disk.EnableUUID**: Specify **TRUE**.
8. In the **Virtual Hardware** panel of the **Customize hardware** tab, modify the specified values as required. Ensure that the amount of RAM, CPU, and disk storage meets the minimum requirements for the machine type. If many networks exist, select **Add New Device** > **Network Adapter**, and then enter your network information in the fields provided by the **New Network** menu item.
9. Complete the remaining configuration steps. On clicking the **Finish** button, you have completed the cloning operation.

- From the **Virtual Machines** tab, right-click on your VM and then select **Power → Power On**.

Next steps


- Continue to create more compute machines for your cluster.

23.3.5.11. Disk partitioning

In most cases, data partitions are originally created by installing RHCOS, rather than by installing another operating system. In such cases, the OpenShift Container Platform installer should be allowed to configure your disk partitions.

However, there are two cases where you might want to intervene to override the default partitioning when installing an OpenShift Container Platform node:

- Create separate partitions:** For greenfield installations on an empty disk, you might want to add separate storage to a partition. This is officially supported for making **/var** or a subdirectory of **/var**, such as **/var/lib/etcd**, a separate partition, but not both.



IMPORTANT

For disk sizes larger than 100GB, and especially disk sizes larger than 1TB, create a separate **/var** partition. See "Creating a separate **/var** partition" and this [Red Hat Knowledgebase article](#) for more information.



IMPORTANT

Kubernetes supports only two file system partitions. If you add more than one partition to the original configuration, Kubernetes cannot monitor all of them.


- Retain existing partitions:** For a brownfield installation where you are reinstalling OpenShift Container Platform on an existing node and want to retain data partitions installed from your previous operating system, there are both boot arguments and options to **coreos-installer** that allow you to retain existing data partitions.

Creating a separate **/var** partition

In general, disk partitioning for OpenShift Container Platform should be left to the installer. However, there are cases where you might want to create separate partitions in a part of the filesystem that you expect to grow.

OpenShift Container Platform supports the addition of a single partition to attach storage to either the **/var** partition or a subdirectory of **/var**. For example:

- /var/lib/containers:** Holds container-related content that can grow as more images and containers are added to a system.
- /var/lib/etcd:** Holds data that you might want to keep separate for purposes such as performance optimization of etcd storage.
- /var:** Holds data that you might want to keep separate for purposes such as auditing.



IMPORTANT

For disk sizes larger than 100GB, and especially larger than 1TB, create a separate **/var** partition.

Storing the contents of a **/var** directory separately makes it easier to grow storage for those areas as needed and reinstall OpenShift Container Platform at a later date and keep that data intact. With this method, you will not have to pull all your containers again, nor will you have to copy massive log files when you update systems.

Because **/var** must be in place before a fresh installation of Red Hat Enterprise Linux CoreOS (RHCOS), the following procedure sets up the separate **/var** partition by creating a machine config manifest that is inserted during the **openshift-install** preparation phases of an OpenShift Container Platform installation.

Procedure

1. Create a directory to hold the OpenShift Container Platform installation files:

```
$ mkdir $HOME/clusterconfig
```

2. Run **openshift-install** to create a set of files in the **manifest** and **openshift** subdirectories. Answer the system questions as you are prompted:

```
$ openshift-install create manifests --dir $HOME/clusterconfig
? SSH Public Key ...
$ ls $HOME/clusterconfig/openshift/
99_kubeadmin-password-secret.yaml
99_openshift-cluster-api_master-machines-0.yaml
99_openshift-cluster-api_master-machines-1.yaml
99_openshift-cluster-api_master-machines-2.yaml
...
```

3. Create a Butane config that configures the additional partition. For example, name the file **\$HOME/clusterconfig/98-var-partition.bu**, change the disk device name to the name of the storage device on the **worker** systems, and set the storage size as appropriate. This example places the **/var** directory on a separate partition:

```
variant: openshift
version: 4.16.0
metadata:
  labels:
    machineconfiguration.openshift.io/role: worker
  name: 98-var-partition
storage:
  disks:
    - device: /dev/disk/by-id/<device_name> 1
      partitions:
        - label: var
          start_mib: <partition_start_offset> 2
          size_mib: <partition_size> 3
          number: 5
      filesystems:
        - device: /dev/disk/by-partlabel/var
          path: /var
          format: xfs
          mount_options: [defaults, prjquota] 4
          with_mount_unit: true
```

- 1 The storage device name of the disk that you want to partition.
- 2 When adding a data partition to the boot disk, a minimum value of 25000 mebibytes is recommended. The root file system is automatically resized to fill all available space up to the specified offset. If no value is specified, or if the specified value is smaller than the recommended minimum, the resulting root file system will be too small, and future reinstalls of RHCOS might overwrite the beginning of the data partition.
- 3 The size of the data partition in mebibytes.
- 4 The **prjquota** mount option must be enabled for filesystems used for container storage.



NOTE

When creating a separate **/var** partition, you cannot use different instance types for worker nodes, if the different instance types do not have the same device name.

4. Create a manifest from the Butane config and save it to the **clusterconfig/openshift** directory. For example, run the following command:

```
$ butane $HOME/clusterconfig/98-var-partition.bu -o $HOME/clusterconfig/openshift/98-var-partition.yaml
```

5. Run **openshift-install** again to create Ignition configs from a set of files in the **manifest** and **openshift** subdirectories:

```
$ openshift-install create ignition-configs --dir $HOME/clusterconfig
$ ls $HOME/clusterconfig/
auth bootstrap.ign master.ign metadata.json worker.ign
```

Now you can use the Ignition config files as input to the vSphere installation procedures to install Red Hat Enterprise Linux CoreOS (RHCOS) systems.

23.3.5.12. Waiting for the bootstrap process to complete

The OpenShift Container Platform bootstrap process begins after the cluster nodes first boot into the persistent RHCOS environment that has been installed to disk. The configuration information provided through the Ignition config files is used to initialize the bootstrap process and install OpenShift Container Platform on the machines. You must wait for the bootstrap process to complete.

Prerequisites

- You have created the Ignition config files for your cluster.
- You have configured suitable network, DNS and load balancing infrastructure.
- You have obtained the installation program and generated the Ignition config files for your cluster.
- You installed RHCOS on your cluster machines and provided the Ignition config files that the OpenShift Container Platform installation program generated.

Procedure

1. Monitor the bootstrap process:

```
$ ./openshift-install --dir <installation_directory> wait-for bootstrap-complete \ 1
--log-level=info 2
```


- 1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.
- 2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Example output

```
INFO Waiting up to 30m0s for the Kubernetes API at https://api.test.example.com:6443...
INFO API v1.29.4 up
INFO Waiting up to 30m0s for bootstrapping to complete...
INFO It is now safe to remove the bootstrap resources
```

The command succeeds when the Kubernetes API server signals that it has been bootstrapped on the control plane machines.

2. After the bootstrap process is complete, remove the bootstrap machine from the load balancer.



IMPORTANT

You must remove the bootstrap machine from the load balancer at this point. You can also remove or reformat the bootstrap machine itself.

23.3.5.13. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.
- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

23.3.5.14. Approving the certificate signing requests for your machines

When you add machines to a cluster, two pending certificate signing requests (CSRs) are generated for each machine that you added. You must confirm that these CSRs are approved or, if necessary, approve them yourself. The client requests must be approved first, followed by the server requests.

Prerequisites

- You added machines to your cluster.

Procedure

- Confirm that the cluster recognizes the machines:

```
$ oc get nodes
```

Example output

```
NAME      STATUS  ROLES  AGE  VERSION
master-0  Ready   master 63m  v1.29.4
master-1  Ready   master 63m  v1.29.4
master-2  Ready   master 64m  v1.29.4
```

The output lists all of the machines that you created.



NOTE

The preceding output might not include the compute nodes, also known as worker nodes, until some CSRs are approved.

- Review the pending CSRs and ensure that you see the client requests with the **Pending** or **Approved** status for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

```
NAME      AGE  REQUESTOR                                     CONDITION
csr-8b2br 15m  system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper  Pending
csr-8vnps 15m  system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper  Pending
...
```



In this example, two machines are joining the cluster. You might see more approved CSRs in the list.

- If the CSRs were not approved, after all of the pending CSRs for the machines you added are in **Pending** status, approve the CSRs for your cluster machines:



NOTE

Because the CSRs rotate automatically, approve your CSRs within an hour of adding the machines to the cluster. If you do not approve them within an hour, the certificates will rotate, and more than two certificates will be present for each node. You must approve all of these certificates. After the client CSR is approved, the Kubelet creates a secondary CSR for the serving certificate, which requires manual approval. Then, subsequent serving certificate renewal requests are automatically approved by the **machine-approver** if the Kubelet requests a new certificate with identical parameters.



NOTE

For clusters running on platforms that are not machine API enabled, such as bare metal and other user-provisioned infrastructure, you must implement a method of automatically approving the kubelet serving certificate requests (CSRs). If a request is not approved, then the **oc exec**, **oc rsh**, and **oc logs** commands cannot succeed, because a serving certificate is required when the API server connects to the kubelet. Any operation that contacts the Kubelet endpoint requires this certificate approval to be in place. The method must watch for new CSRs, confirm that the CSR was submitted by the **node-bootstrap** service account in the **system:node** or **system:admin** groups, and confirm the identity of the node.

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{"\n"}\n{{end}}{{end}}' | xargs --no-run-if-empty oc adm certificate approve
```



NOTE

Some Operators might not become available until some CSRs are approved.

- Now that your client requests are approved, you must review the server requests for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

NAME	AGE	REQUESTOR	CONDITION
csr-bfd72	5m26s	system:node:ip-10-0-50-126.us-east-2.compute.internal	Pending
csr-c57lv	5m26s	system:node:ip-10-0-95-157.us-east-2.compute.internal	Pending
...			

5. If the remaining CSRs are not approved, and are in the **Pending** status, approve the CSRs for your cluster machines:

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

1 **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{"\n"}{{end}}{{end}}' | xargs oc adm certificate approve
```

6. After all client and server CSRs have been approved, the machines have the **Ready** status. Verify this by running the following command:

```
$ oc get nodes
```

Example output

NAME	STATUS	ROLES	AGE	VERSION
master-0	Ready	master	73m	v1.29.4
master-1	Ready	master	73m	v1.29.4
master-2	Ready	master	74m	v1.29.4
worker-0	Ready	worker	11m	v1.29.4
worker-1	Ready	worker	11m	v1.29.4



NOTE

It can take a few minutes after approval of the server CSRs for the machines to transition to the **Ready** status.

Additional information

- For more information on CSRs, see [Certificate Signing Requests](#).

23.3.5.15. Initial Operator configuration

After the control plane initializes, you must immediately configure some Operators so that they all become available.

Prerequisites

- Your control plane has initialized.

Procedure

1. Watch the cluster components come online:

```
$ watch -n5 oc get clusteroperators
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.16.0	True	False	False	19m
baremetal	4.16.0	True	False	False	37m
cloud-credential	4.16.0	True	False	False	40m
cluster-autoscaler	4.16.0	True	False	False	37m
config-operator	4.16.0	True	False	False	38m
console	4.16.0	True	False	False	26m
csi-snapshot-controller	4.16.0	True	False	False	37m
dns	4.16.0	True	False	False	37m
etcd	4.16.0	True	False	False	36m
image-registry	4.16.0	True	False	False	31m
ingress	4.16.0	True	False	False	30m
insights	4.16.0	True	False	False	31m
kube-apiserver	4.16.0	True	False	False	26m
kube-controller-manager	4.16.0	True	False	False	36m
kube-scheduler	4.16.0	True	False	False	36m
kube-storage-version-migrator	4.16.0	True	False	False	37m
machine-api	4.16.0	True	False	False	29m
machine-approver	4.16.0	True	False	False	37m
machine-config	4.16.0	True	False	False	36m
marketplace	4.16.0	True	False	False	37m
monitoring	4.16.0	True	False	False	29m
network	4.16.0	True	False	False	38m
node-tuning	4.16.0	True	False	False	37m
openshift-apiserver	4.16.0	True	False	False	32m
openshift-controller-manager	4.16.0	True	False	False	30m
openshift-samples	4.16.0	True	False	False	32m
operator-lifecycle-manager	4.16.0	True	False	False	37m
operator-lifecycle-manager-catalog	4.16.0	True	False	False	37m
operator-lifecycle-manager-packageserver	4.16.0	True	False	False	32m
service-ca	4.16.0	True	False	False	38m
storage	4.16.0	True	False	False	37m

2. Configure the Operators that are not available.

23.3.5.15.1. Disabling the default OperatorHub catalog sources

Operator catalogs that source content provided by Red Hat and community projects are configured for OperatorHub by default during an OpenShift Container Platform installation. In a restricted network environment, you must disable the default catalogs as a cluster administrator.

Procedure

- Disable the sources for the default catalogs by adding **disableAllDefaultSources: true** to the **OperatorHub** object:

```
$ oc patch OperatorHub cluster --type json \
  -p '[{"op": "add", "path": "/spec/disableAllDefaultSources", "value": true}]'
```

TIP

Alternatively, you can use the web console to manage catalog sources. From the **Administration** → **Cluster Settings** → **Configuration** → **OperatorHub** page, click the **Sources** tab, where you can create, update, delete, disable, and enable individual sources.

23.3.5.15.2. Image registry storage configuration

The Image Registry Operator is not initially available for platforms that do not provide default storage. After installation, you must configure your registry to use storage so that the Registry Operator is made available.

Instructions are shown for configuring a persistent volume, which is required for production clusters. Where applicable, instructions are shown for configuring an empty directory as the storage location, which is available for only non-production clusters.


Additional instructions are provided for allowing the image registry to use block storage types by using the **Recreate** rollout strategy during upgrades.

23.3.5.15.2.1. Configuring registry storage for VMware vSphere

As a cluster administrator, following installation you must configure your registry to use storage.

Prerequisites


- Cluster administrator permissions.
- A cluster on VMware vSphere.
- Persistent storage provisioned for your cluster, such as Red Hat OpenShift Data Foundation.



IMPORTANT

OpenShift Container Platform supports **ReadWriteOnce** access for image registry storage when you have only one replica. **ReadWriteOnce** access also requires that the registry uses the **Recreate** rollout strategy. To deploy an image registry that supports high availability with two or more replicas, **ReadWriteMany** access is required.

- Must have "100Gi" capacity.




IMPORTANT

Testing shows issues with using the NFS server on RHEL as storage backend for core services. This includes the OpenShift Container Registry and Quay, Prometheus for monitoring storage, and Elasticsearch for logging storage. Therefore, using RHEL NFS to back PVs used by core services is not recommended.

Other NFS implementations on the marketplace might not have these issues. Contact the individual NFS implementation vendor for more information on any testing that was possibly completed against these OpenShift Container Platform core components.

Procedure

1. To configure your registry to use storage, change the **spec.storage.pvc** in the **configs.imageregistry/cluster** resource.



NOTE

When you use shared storage, review your security settings to prevent outside access.

2. Verify that you do not have a registry pod:

```
$ oc get pod -n openshift-image-registry -l docker-registry=default
```

Example output

```
No resources found in openshift-image-registry namespace
```



NOTE

If you do have a registry pod in your output, you do not need to continue with this procedure.

3. Check the registry configuration:

```
$ oc edit configs.imageregistry.operator.openshift.io
```

Example output

```
storage:
  pvc:
    claim: 1
```

- 1** Leave the **claim** field blank to allow the automatic creation of an **image-registry-storage** persistent volume claim (PVC). The PVC is generated based on the default storage class. However, be aware that the default storage class might provide ReadWriteOnce (RWO) volumes, such as a RADOS Block Device (RBD), which can cause issues when you replicate to more than one replica.

4. Check the **clusteroperator** status:

```
$ oc get clusteroperator image-registry
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED
image-registry	4.7	True	False	False

23.3.5.15.2.2. Configuring storage for the image registry in non-production clusters

You must configure storage for the Image Registry Operator. For non-production clusters, you can set the image registry to an empty directory. If you do so, all images are lost if you restart the registry.

Procedure

- To set the image registry storage to an empty directory:

```
$ oc patch configs.imageregistry.operator.openshift.io cluster --type merge --patch '{"spec": {"storage":{"emptyDir":{}}}'
```



WARNING

Configure this option for only non-production clusters.

If you run this command before the Image Registry Operator initializes its components, the **oc patch** command fails with the following error:

```
Error from server (NotFound): configs.imageregistry.operator.openshift.io "cluster" not found
```

Wait a few minutes and run the command again.

23.3.5.15.2.3. Configuring block registry storage for VMware vSphere

To allow the image registry to use block storage types such as vSphere Virtual Machine Disk (VMDK) during upgrades as a cluster administrator, you can use the **Recreate** rollout strategy.



IMPORTANT

Block storage volumes are supported but not recommended for use with image registry on production clusters. An installation where the registry is configured on block storage is not highly available because the registry cannot have more than one replica.

Procedure

- Enter the following command to set the image registry storage as a block storage type, patch the registry so that it uses the **Recreate** rollout strategy, and runs with only **1** replica:

```
$ oc patch config.imageregistry.operator.openshift.io/cluster --type=merge -p '{"spec": {"rolloutStrategy":"Recreate","replicas":1}}'
```

- Provision the PV for the block storage device, and create a PVC for that volume. The requested block volume uses the ReadWriteOnce (RWO) access mode.
 - Create a **pvc.yaml** file with the following contents to define a VMware vSphere **PersistentVolumeClaim** object:

```
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
```

```

name: image-registry-storage 1
namespace: openshift-image-registry 2
spec:
  accessModes:
  - ReadWriteOnce 3
  resources:
  requests:
    storage: 100Gi 4

```

- 1** A unique name that represents the **PersistentVolumeClaim** object.
- 2** The namespace for the **PersistentVolumeClaim** object, which is **openshift-image-registry**.
- 3** The access mode of the persistent volume claim. With **ReadWriteOnce**, the volume can be mounted with read and write permissions by a single node.
- 4** The size of the persistent volume claim.

b. Enter the following command to create the **PersistentVolumeClaim** object from the file:

```
$ oc create -f pvc.yaml -n openshift-image-registry
```

3. Enter the following command to edit the registry configuration so that it references the correct PVC:

```
$ oc edit config.imageregistry.operator.openshift.io -o yaml
```

Example output

```

storage:
  pvc:
    claim: 1

```

- 1** By creating a custom PVC, you can leave the **claim** field blank for the default automatic creation of an **image-registry-storage** PVC.

For instructions about configuring registry storage so that it references the correct PVC, see [Configuring the registry for vSphere](#).

23.3.5.16. Completing installation on user-provisioned infrastructure

After you complete the Operator configuration, you can finish installing the cluster on infrastructure that you provide.

Prerequisites

- Your control plane has initialized.
- You have completed the initial Operator configuration.

Procedure

1. Confirm that all the cluster components are online with the following command:

```
$ watch -n5 oc get clusteroperators
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.16.0	True	False	False	19m
baremetal	4.16.0	True	False	False	37m
cloud-credential	4.16.0	True	False	False	40m
cluster-autoscaler	4.16.0	True	False	False	37m
config-operator	4.16.0	True	False	False	38m
console	4.16.0	True	False	False	26m
csi-snapshot-controller	4.16.0	True	False	False	37m
dns	4.16.0	True	False	False	37m
etcd	4.16.0	True	False	False	36m
image-registry	4.16.0	True	False	False	31m
ingress	4.16.0	True	False	False	30m
insights	4.16.0	True	False	False	31m
kube-apiserver	4.16.0	True	False	False	26m
kube-controller-manager	4.16.0	True	False	False	36m
kube-scheduler	4.16.0	True	False	False	36m
kube-storage-version-migrator	4.16.0	True	False	False	37m
machine-api	4.16.0	True	False	False	29m
machine-approver	4.16.0	True	False	False	37m
machine-config	4.16.0	True	False	False	36m
marketplace	4.16.0	True	False	False	37m
monitoring	4.16.0	True	False	False	29m
network	4.16.0	True	False	False	38m
node-tuning	4.16.0	True	False	False	37m
openshift-apiserver	4.16.0	True	False	False	32m
openshift-controller-manager	4.16.0	True	False	False	30m
openshift-samples	4.16.0	True	False	False	32m
operator-lifecycle-manager	4.16.0	True	False	False	37m
operator-lifecycle-manager-catalog	4.16.0	True	False	False	37m
operator-lifecycle-manager-packageserver	4.16.0	True	False	False	32m
service-ca	4.16.0	True	False	False	38m
storage	4.16.0	True	False	False	37m

Alternatively, the following command notifies you when all of the clusters are available. It also retrieves and displays credentials:

```
$ ./openshift-install --dir <installation_directory> wait-for install-complete 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

Example output

```
INFO Waiting up to 30m0s for the cluster to initialize...
```


The command succeeds when the Cluster Version Operator finishes deploying the OpenShift Container Platform cluster from Kubernetes API server.



IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

2. Confirm that the Kubernetes API server is communicating with the pods.

a. To view a list of all pods, use the following command:

```
$ oc get pods --all-namespaces
```

Example output

```

NAMESPACE           NAME                                     READY  STATUS   RESTARTS  AGE
openshift-apiserver-operator  openshift-apiserver-operator-85cb746d55-zqhs8  1/1    Running   1          9m
openshift-apiserver          apiserver-67b9g                                1/1    Running   0          3m
openshift-apiserver          apiserver-ljcmx                                1/1    Running   0          1m
openshift-apiserver          apiserver-z25h4                                1/1    Running   0          2m
openshift-authentication-operator  authentication-operator-69d5d8bf84-vh2n8      1/1    Running   0          5m
...

```

b. View the logs for a pod that is listed in the output of the previous command by using the following command:

```
$ oc logs <pod_name> -n <namespace> 1
```

- 1** Specify the pod name and namespace, as shown in the output of the previous command.

If the pod logs display, the Kubernetes API server can communicate with the cluster machines.

3. For an installation with Fibre Channel Protocol (FCP), additional steps are required to enable multipathing. Do not enable multipathing during installation. See "Enabling multipathing with kernel arguments on RHCOS" in the *Postinstallation machine configuration tasks* documentation for more information.
4. Register your cluster on the [Cluster registration](#) page.

You can add extra compute machines after the cluster installation is completed by following [Adding compute machines to vSphere](#).

23.3.5.17. Configuring vSphere DRS anti-affinity rules for control plane nodes

vSphere Distributed Resource Scheduler (DRS) anti-affinity rules can be configured to support higher availability of OpenShift Container Platform Control Plane nodes. Anti-affinity rules ensure that the vSphere Virtual Machines for the OpenShift Container Platform Control Plane nodes are not scheduled to the same vSphere Host.



IMPORTANT

- The following information applies to compute DRS only and does not apply to storage DRS.
- The **govc** command is an open-source command available from VMware; it is not available from Red Hat. The **govc** command is not supported by the Red Hat support.
- Instructions for downloading and installing **govc** are found on the VMware documentation website.

Create an anti-affinity rule by running the following command:

Example command

```
$ govc cluster.rule.create \
-name openshift4-control-plane-group \
-dc MyDatacenter -cluster MyCluster \
-enable \
-anti-affinity master-0 master-1 master-2
```

After creating the rule, your control plane nodes are automatically migrated by vSphere so they are not running on the same hosts. This might take some time while vSphere reconciles the new rule. Successful command completion is shown in the following procedure.



NOTE

The migration occurs automatically and might cause brief OpenShift API outage or latency until the migration finishes.

The vSphere DRS anti-affinity rules need to be updated manually in the event of a control plane VM name change or migration to a new vSphere Cluster.

Procedure

1. Remove any existing DRS anti-affinity rule by running the following command:

```
$ govc cluster.rule.remove \
-name openshift4-control-plane-group \
-dc MyDatacenter -cluster MyCluster
```

Example Output

```
[13-10-22 09:33:24] Reconfigure /MyDatacenter/host/MyCluster...OK
```

2. Create the rule again with updated names by running the following command:

```
$ govc cluster.rule.create \
-name openshift4-control-plane-group \
-dc MyDatacenter -cluster MyOtherCluster \
-enable \
-anti-affinity master-0 master-1 master-2
```

23.3.5.18. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service

23.3.5.19. Next steps

- [Customize your cluster](#).
- If the mirror registry that you used to install your cluster has a trusted CA, add it to the cluster by [configuring additional trust stores](#).
- If necessary, you can [opt out of remote health reporting](#) .
- Optional: [View the events from the vSphere Problem Detector Operator](#) to determine if the cluster has permission or storage configuration issues.
- Optional: if you created encrypted virtual machines, [create an encrypted storage class](#) .

23.4. INSTALLING A CLUSTER ON VSPHERE USING THE ASSISTED INSTALLER

You can install OpenShift Container Platform on on-premise hardware or on-premise VMs by using the Assisted Installer. Installing OpenShift Container Platform by using the Assisted Installer supports **x86_64**, **AArch64**, **ppc64le**, and **s390x** CPU architectures.

The Assisted Installer is a user-friendly installation solution offered on the Red Hat Hybrid Cloud Console.

23.4.1. Additional resources

- [Installing OpenShift Container Platform with the Assisted Installer](#)

23.5. INSTALLING A CLUSTER ON VSPHERE USING THE AGENT-BASED INSTALLER

The Agent-based installation method provides the flexibility to boot your on-premises servers in any way that you choose. It combines the ease of use of the Assisted Installation service with the ability to run offline, including in air-gapped environments.

Agent-based installation is a subcommand of the OpenShift Container Platform installer. It generates a bootable ISO image containing all of the information required to deploy an OpenShift Container Platform cluster with an available release image.

23.5.1. Additional resources

- [Preparing to install with the Agent-based Installer](#)

23.6. INSTALLING A THREE-NODE CLUSTER ON VSPHERE

In OpenShift Container Platform version 4.16, you can install a three-node cluster on VMware vSphere. A three-node cluster consists of three control plane machines, which also act as compute machines. This type of cluster provides a smaller, more resource efficient cluster, for cluster administrators and developers to use for testing, development, and production.

You can install a three-node cluster using either installer-provisioned or user-provisioned infrastructure.

23.6.1. Configuring a three-node cluster

You configure a three-node cluster by setting the number of worker nodes to **0** in the **install-config.yaml** file before deploying the cluster. Setting the number of worker nodes to **0** ensures that the control plane machines are schedulable. This allows application workloads to be scheduled to run from the control plane nodes.



NOTE

Because application workloads run from control plane nodes, additional subscriptions are required, as the control plane nodes are considered to be compute nodes.

Prerequisites

- You have an existing **install-config.yaml** file.

Procedure

1. Set the number of compute replicas to **0** in your **install-config.yaml** file, as shown in the following **compute** stanza:

Example **install-config.yaml** file for a three-node cluster

```
apiVersion: v1
baseDomain: example.com
```

```

compute:
- name: worker
  platform: {}
  replicas: 0
# ...

```

2. If you are deploying a cluster with user-provisioned infrastructure:

- Configure your application ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes. In a three-node cluster, the Ingress Controller pods run on the control plane nodes. For more information, see the "Load balancing requirements for user-provisioned infrastructure".
- After you create the Kubernetes manifest files, make sure that the **spec.mastersSchedulable** parameter is set to **true** in **cluster-scheduler-02-config.yml** file. You can locate this file in **<installation_directory>/manifests**. For more information, see "Creating the Kubernetes manifest and Ignition config files" in "Installing a cluster on vSphere with user-provisioned infrastructure".
- Do not create additional worker nodes.

Example cluster-scheduler-02-config.yml file for a three-node cluster

```

apiVersion: config.openshift.io/v1
kind: Scheduler
metadata:
  creationTimestamp: null
  name: cluster
spec:
  mastersSchedulable: true
  policy:
    name: ""
status: {}

```

23.6.2. Next steps

- [Installing a cluster on vSphere with customizations](#)
- [Installing a cluster on vSphere with user-provisioned infrastructure](#)

23.7. UNINSTALLING A CLUSTER ON VSPHERE THAT USES INSTALLER-PROVISIONED INFRASTRUCTURE

You can remove a cluster that you deployed in your VMware vSphere instance by using installer-provisioned infrastructure.



NOTE

When you run the **openshift-install destroy cluster** command to uninstall OpenShift Container Platform, vSphere volumes are not automatically deleted. The cluster administrator must manually find the vSphere volumes and delete them.

23.7.1. Removing a cluster that uses installer-provisioned infrastructure

You can remove a cluster that uses installer-provisioned infrastructure from your cloud.



NOTE

After uninstallation, check your cloud provider for any resources not removed properly, especially with User Provisioned Infrastructure (UPI) clusters. There might be resources that the installer did not create or that the installer is unable to access.

Prerequisites

- You have a copy of the installation program that you used to deploy the cluster.
- You have the files that the installation program generated when you created your cluster.

Procedure

1. From the directory that contains the installation program on the computer that you used to install the cluster, run the following command:

```
$. /openshift-install destroy cluster \  
--dir <installation_directory> --log-level info 1 2
```

1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2 To view different details, specify **warn**, **debug**, or **error** instead of **info**.



NOTE

You must specify the directory that contains the cluster definition files for your cluster. The installation program requires the **metadata.json** file in this directory to delete the cluster.

2. Optional: Delete the **<installation_directory>** directory and the OpenShift Container Platform installation program.

23.8. USING THE VSPHERE PROBLEM DETECTOR OPERATOR

23.8.1. About the vSphere Problem Detector Operator

The vSphere Problem Detector Operator checks clusters that are deployed on vSphere for common installation and misconfiguration issues that are related to storage.

The Operator runs in the **openshift-cluster-storage-operator** namespace and is started by the Cluster Storage Operator when the Cluster Storage Operator detects that the cluster is deployed on vSphere. The vSphere Problem Detector Operator communicates with the vSphere vCenter Server to determine the virtual machines in the cluster, the default datastore, and other information about the vSphere vCenter Server configuration. The Operator uses the credentials from the Cloud Credential Operator to connect to vSphere.

The Operator runs the checks according to the following schedule:

- The checks run every 8 hours.
- If any check fails, the Operator runs the checks again in intervals of 1 minute, 2 minutes, 4, 8, and so on. The Operator doubles the interval up to a maximum interval of 8 hours.
- When all checks pass, the schedule returns to an 8 hour interval.

The Operator increases the frequency of the checks after a failure so that the Operator can report success quickly after the failure condition is remedied. You can run the Operator manually for immediate troubleshooting information.

23.8.2. Running the vSphere Problem Detector Operator checks

You can override the schedule for running the vSphere Problem Detector Operator checks and run the checks immediately.

The vSphere Problem Detector Operator automatically runs the checks every 8 hours. However, when the Operator starts, it runs the checks immediately. The Operator is started by the Cluster Storage Operator when the Cluster Storage Operator starts and determines that the cluster is running on vSphere. To run the checks immediately, you can scale the vSphere Problem Detector Operator to **0** and back to **1** so that it restarts the vSphere Problem Detector Operator.

Prerequisites

- Access to the cluster as a user with the **cluster-admin** role.

Procedure

1. Scale the Operator to **0**:

```
$ oc scale deployment/vsphere-problem-detector-operator --replicas=0 \
-n openshift-cluster-storage-operator
```

If the deployment does not scale to zero immediately, you can run the following command to wait for the pods to exit:

```
$ oc wait pods -l name=vsphere-problem-detector-operator \
--for=delete --timeout=5m -n openshift-cluster-storage-operator
```

2. Scale the Operator back to **1**:

```
$ oc scale deployment/vsphere-problem-detector-operator --replicas=1 \
-n openshift-cluster-storage-operator
```

3. Delete the old leader lock to speed up the new leader election for the Cluster Storage Operator:

```
$ oc delete -n openshift-cluster-storage-operator \
cm vsphere-problem-detector-lock
```

Verification

- View the events or logs that are generated by the vSphere Problem Detector Operator. Confirm that the events or logs have recent timestamps.

23.8.3. Viewing the events from the vSphere Problem Detector Operator

After the vSphere Problem Detector Operator runs and performs the configuration checks, it creates events that can be viewed from the command line or from the OpenShift Container Platform web console.

Procedure

- To view the events by using the command line, run the following command:

```
$ oc get event -n openshift-cluster-storage-operator \
  --sort-by={.metadata.creationTimestamp}
```

Example output

```
16m Normal Started pod/vsphere-problem-detector-operator-xxxxx Started
container vsphere-problem-detector
16m Normal Created pod/vsphere-problem-detector-operator-xxxxx Created
container vsphere-problem-detector
16m Normal LeaderElection configmap/vsphere-problem-detector-lock vsphere-
problem-detector-operator-xxxxx became leader
```

- To view the events by using the OpenShift Container Platform web console, navigate to **Home** → **Events** and select **openshift-cluster-storage-operator** from the **Project** menu.

23.8.4. Viewing the logs from the vSphere Problem Detector Operator

After the vSphere Problem Detector Operator runs and performs the configuration checks, it creates log records that can be viewed from the command line or from the OpenShift Container Platform web console.

Procedure

- To view the logs by using the command line, run the following command:

```
$ oc logs deployment/vsphere-problem-detector-operator \
  -n openshift-cluster-storage-operator
```

Example output

```
I0108 08:32:28.445696 1 operator.go:209] ClusterInfo passed
I0108 08:32:28.451029 1 datastore.go:57] CheckStorageClasses checked 1 storage
classes, 0 problems found
I0108 08:32:28.451047 1 operator.go:209] CheckStorageClasses passed
I0108 08:32:28.452160 1 operator.go:209] CheckDefaultDatastore passed
I0108 08:32:28.480648 1 operator.go:271] CheckNodeDiskUUID:<host_name> passed
I0108 08:32:28.480685 1 operator.go:271] CheckNodeProviderID:<host_name> passed
```

- To view the Operator logs with the OpenShift Container Platform web console, perform the following steps:
 - Navigate to **Workloads** → **Pods**.
 - Select **openshift-cluster-storage-operator** from the **Projects** menu.

- c. Click the link for the **vsphere-problem-detector-operator** pod.
- d. Click the **Logs** tab on the **Pod details** page to view the logs.

23.8.5. Configuration checks run by the vSphere Problem Detector Operator

The following tables identify the configuration checks that the vSphere Problem Detector Operator runs. Some checks verify the configuration of the cluster. Other checks verify the configuration of each node in the cluster.

Table 23.40. Cluster configuration checks

Name	Description
CheckDefaultDatastore	<p>Verifies that the default datastore name in the vSphere configuration is short enough for use with dynamic provisioning.</p> <p>If this check fails, you can expect the following:</p> <ul style="list-style-type: none"> ● systemd logs errors to the journal such as Failed to set up mount unit: Invalid argument. ● systemd does not unmount volumes if the virtual machine is shut down or rebooted without draining all the pods from the node. <p>If this check fails, reconfigure vSphere with a shorter name for the default datastore.</p>
CheckFolderPermissions	<p>Verifies the permission to list volumes in the default datastore. This permission is required to create volumes. The Operator verifies the permission by listing the / and /kubevols directories. The root directory must exist. It is acceptable if the /kubevols directory does not exist when the check runs. The /kubevols directory is created when the datastore is used with dynamic provisioning if the directory does not already exist.</p> <p>If this check fails, review the required permissions for the vCenter account that was specified during the OpenShift Container Platform installation.</p>
CheckStorageClasses	<p>Verifies the following:</p> <ul style="list-style-type: none"> ● The fully qualified path to each persistent volume that is provisioned by this storage class is less than 255 characters. ● If a storage class uses a storage policy, the storage class must use one policy only and that policy must be defined.
CheckTaskPermissions	<p>Verifies the permission to list recent tasks and datastores.</p>
ClusterInfo	<p>Collects the cluster version and UUID from vSphere vCenter.</p>

Table 23.41. Node configuration checks

Name	Description
CheckNodeDiskUUID	<p>Verifies that all the vSphere virtual machines are configured with disk.enableUUID=TRUE.</p> <p>If this check fails, see the How to check 'disk.EnableUUID' parameter from VM in vSphere Red Hat Knowledgebase solution.</p>
CheckNodeProviderID	<p>Verifies that all nodes are configured with the ProviderID from vSphere vCenter. This check fails when the output from the following command does not include a provider ID for each node.</p> <pre>\$ oc get nodes -o custom-columns=NAME:.metadata.name,PROVIDER_ID:.spec.providerID,UUID:.status.nodeInfo.systemUUID</pre> <p>If this check fails, refer to the vSphere product documentation for information about setting the provider ID for each node in the cluster.</p>
CollectNodeESXiVersion	Reports the version of the ESXi hosts that run nodes.
CollectNodeHWVersion	Reports the virtual machine hardware version for a node.

23.8.6. About the storage class configuration check

The names for persistent volumes that use vSphere storage are related to the datastore name and cluster ID.

When a persistent volume is created, **systemd** creates a mount unit for the persistent volume. The **systemd** process has a 255 character limit for the length of the fully qualified path to the VDMK file that is used for the persistent volume.

The fully qualified path is based on the naming conventions for **systemd** and vSphere. The naming conventions use the following pattern:

```
/var/lib/kubelet/plugins/kubernetes.io/vsphere-volume/mounts/[-<datastore>] 00000000-0000-0000-0000-0000000000000000/<cluster_id>-dynamic-pvc-00000000-0000-0000-0000-000000000000.vmdk
```

- The naming conventions require 205 characters of the 255 character limit.
- The datastore name and the cluster ID are determined from the deployment.
- The datastore name and cluster ID are substituted into the preceding pattern. Then the path is processed with the **systemd-escape** command to escape special characters. For example, a hyphen character uses four characters after it is escaped. The escaped value is **\x2d**.
- After processing with **systemd-escape** to ensure that **systemd** can access the fully qualified path to the VDMK file, the length of the path must be less than 255 characters.

23.8.7. Metrics for the vSphere Problem Detector Operator

The vSphere Problem Detector Operator exposes the following metrics for use by the OpenShift Container Platform monitoring stack.

Table 23.42. Metrics exposed by the vSphere Problem Detector Operator

Name	Description
vsphere_cluster_check_total	Cumulative number of cluster-level checks that the vSphere Problem Detector Operator performed. This count includes both successes and failures.
vsphere_cluster_check_errors	Number of failed cluster-level checks that the vSphere Problem Detector Operator performed. For example, a value of 1 indicates that one cluster-level check failed.
vsphere_esxi_version_total	Number of ESXi hosts with a specific version. Be aware that if a host runs more than one node, the host is counted only once.
vsphere_node_check_total	Cumulative number of node-level checks that the vSphere Problem Detector Operator performed. This count includes both successes and failures.
vsphere_node_check_errors	Number of failed node-level checks that the vSphere Problem Detector Operator performed. For example, a value of 1 indicates that one node-level check failed.
vsphere_node_hw_version_total	Number of vSphere nodes with a specific hardware version.
vsphere_vcenter_info	Information about the vSphere vCenter Server.

23.8.8. Additional resources

- [Monitoring overview](#)

23.9. INSTALLATION CONFIGURATION PARAMETERS FOR VSPHERE

Before you deploy an OpenShift Container Platform cluster on vSphere, you provide parameters to customize your cluster and the platform that hosts it. When you create the **install-config.yaml** file, you provide values for the required parameters through the command line. You can then modify the **install-config.yaml** file to customize your cluster further.

23.9.1. Available installation configuration parameters for vSphere

The following tables specify the required, optional, and vSphere-specific installation configuration parameters that you can set as part of the installation process.



NOTE

After installation, you cannot modify these parameters in the **install-config.yaml** file.

23.9.1.1. Required configuration parameters

Required installation configuration parameters are described in the following table:

Table 23.43. Required parameters

Parameter	Description	Values
<code>apiVersion:</code>	The API version for the install-config.yaml content. The current version is v1 . The installation program may also support older API versions.	String
<code>baseDomain:</code>	The base domain of your cloud provider. The base domain is used to create routes to your OpenShift Container Platform cluster components. The full DNS name for your cluster is a combination of the baseDomain and metadata.name parameter values that uses the <metadata.name> . <baseDomain> format.	A fully-qualified domain or subdomain name, such as example.com .
<code>metadata:</code>	Kubernetes resource ObjectMeta , from which only the name parameter is consumed.	Object
<code>metadata: name:</code>	The name of the cluster. DNS records for the cluster are all subdomains of {{.metadata.name}} . {{.baseDomain}} .	String of lowercase letters and hyphens (-), such as dev .
<code>platform:</code>	The configuration for the specific platform upon which to perform the installation: aws , baremetal , azure , gcp , ibmcloud , nutanix , openstack , powervs , vsphere , or {} . For additional information about platform . <platform> parameters, consult the table for your specific platform that follows.	Object

Parameter	Description	Values
<code>pullSecret:</code>	Get a pull secret from Red Hat OpenShift Cluster Manager to authenticate downloading container images for OpenShift Container Platform components from services such as Quay.io.	<pre>{ "auths":{ "cloud.openshift.com":{ "auth":"b3Blb=", "email":"you@example.com" }, "quay.io":{ "auth":"b3Blb=", "email":"you@example.com" } } }</pre>

23.9.1.2. Network configuration parameters

You can customize your installation configuration based on the requirements of your existing network infrastructure. For example, you can expand the IP address block for the cluster network or provide different IP address blocks than the defaults.

- If you use the Red Hat OpenShift Networking OVN-Kubernetes network plugin, both IPv4 and IPv6 address families are supported.



NOTE

On VMware vSphere, dual-stack networking can specify either IPv4 or IPv6 as the primary address family.

If you configure your cluster to use both IP address families, review the following requirements:


- Both IP families must use the same network interface for the default gateway.
- Both IP families must have the default gateway.
- You must specify IPv4 and IPv6 addresses in the same order for all network configuration parameters. For example, in the following configuration IPv4 addresses are listed before IPv6 addresses.


```
networking:
  clusterNetwork:
    - cidr: 10.128.0.0/14
      hostPrefix: 23
    - cidr: fd00:10:128::/56
      hostPrefix: 64
  serviceNetwork:
    - 172.30.0.0/16
    - fd00:172:16::/112
```

**NOTE**

Globalnet is not supported with Red Hat OpenShift Data Foundation disaster recovery solutions. For regional disaster recovery scenarios, ensure that you use a nonoverlapping range of private IP addresses for the cluster and service networks in each cluster.

Table 23.44. Network parameters

Parameter	Description	Values
<code>networking:</code>	The configuration for the cluster network.	Object  NOTE You cannot modify parameters specified by the networking object after installation.
<code>networking: networkType:</code>	The Red Hat OpenShift Networking network plugin to install.	OVNKubernetes. OVNKubernetes is a CNI plugin for Linux networks and hybrid networks that contain both Linux and Windows servers. The default value is OVNKubernetes .
<code>networking: clusterNetwork:</code>	The IP address blocks for pods. The default value is 10.128.0.0/14 with a host prefix of /23 . If you specify multiple IP address blocks, the blocks must not overlap.	An array of objects. For example: <code>networking: clusterNetwork: - cidr: 10.128.0.0/14 hostPrefix: 23</code>
<code>networking: clusterNetwork: cidr:</code>	Required if you use networking.clusterNetwork . An IP address block. An IPv4 network.	An IP address block in Classless Inter-Domain Routing (CIDR) notation. The prefix length for an IPv4 block is between 0 and 32 .
<code>networking: clusterNetwork: hostPrefix:</code>	The subnet prefix length to assign to each individual node. For example, if hostPrefix is set to 23 then each node is assigned a /23 subnet out of the given cidr . A hostPrefix value of 23 provides 510 ($2^{(32 - 23)} - 2$) pod IP addresses.	A subnet prefix. The default value is 23 .

Parameter	Description	Values
<code>networking: serviceNetwork:</code>	<p>The IP address block for services. The default value is 172.30.0.0/16.</p> <p>The OVN-Kubernetes network plugins supports only a single IP address block for the service network.</p>	<p>An array with an IP address block in CIDR format. For example:</p> <pre>networking: serviceNetwork: - 172.30.0.0/16</pre>
<code>networking: machineNetwork:</code>	<p>The IP address blocks for machines.</p> <p>If you specify multiple IP address blocks, the blocks must not overlap.</p>	<p>An array of objects. For example:</p> <pre>networking: machineNetwork: - cidr: 10.0.0.0/16</pre>
<code>networking: machineNetwork: cidr:</code>	<p>Required if you use networking.machineNetwork. An IP address block. The default value is 10.0.0.0/16 for all platforms other than libvirt and IBM Power® Virtual Server. For libvirt, the default value is 192.168.126.0/24. For IBM Power® Virtual Server, the default value is 192.168.0.0/24.</p>	<p>An IP network block in CIDR notation. For example, 10.0.0.0/16.</p> <div style="display: flex; align-items: center;">  <div> <p>NOTE</p> <p>Set the networking.machineNetwork to match the CIDR that the preferred NIC resides in.</p> </div> </div>

23.9.1.3. Optional configuration parameters

Optional installation configuration parameters are described in the following table:

Table 23.45. Optional parameters

Parameter	Description	Values
<code>additionalTrustBundle:</code>	<p>A PEM-encoded X.509 certificate bundle that is added to the nodes' trusted certificate store. This trust bundle may also be used when a proxy has been configured.</p>	String

Parameter	Description	Values
capabilities:	Controls the installation of optional core cluster components. You can reduce the footprint of your OpenShift Container Platform cluster by disabling optional components. For more information, see the "Cluster capabilities" page in <i>Installing</i> .	String array
capabilities: baselineCapabilitySet:	Selects an initial set of optional capabilities to enable. Valid values are None , v4.11 , v4.12 and vCurrent . The default value is vCurrent .	String
capabilities: additionalEnabledCapabilities:	Extends the set of optional capabilities beyond what you specify in baselineCapabilitySet . You may specify multiple capabilities in this parameter.	String array
cpuPartitioningMode:	Enables workload partitioning, which isolates OpenShift Container Platform services, cluster management workloads, and infrastructure pods to run on a reserved set of CPUs. Workload partitioning can only be enabled during installation and cannot be disabled after installation. While this field enables workload partitioning, it does not configure workloads to use specific CPUs. For more information, see the <i>Workload partitioning</i> page in the <i>Scalability and Performance</i> section.	None or AllNodes . None is the default value.
compute:	The configuration for the machines that comprise the compute nodes.	Array of MachinePool objects.
compute: architecture:	Determines the instruction set architecture of the machines in the pool. Currently, clusters with varied architectures are not supported. All pools must specify the same architecture. Valid values are amd64 (the default).	String

Parameter	Description	Values
<code>compute: name:</code>	Required if you use compute . The name of the machine pool.	worker
<code>compute: platform:</code>	Required if you use compute . Use this parameter to specify the cloud provider to host the worker machines. This parameter value must match the controlPlane.platform parameter value.	aws, azure, gcp, ibmcloud, nutanix, openstack, powervs, vsphere, or {}
<code>compute: replicas:</code>	The number of compute machines, which are also known as worker machines, to provision.	A positive integer greater than or equal to 2 . The default value is 3 .
<code>featureSet:</code>	Enables the cluster for a feature set. A feature set is a collection of OpenShift Container Platform features that are not enabled by default. For more information about enabling a feature set during installation, see "Enabling features using feature gates".	String. The name of the feature set to enable, such as TechPreviewNoUpgrade .
<code>controlPlane:</code>	The configuration for the machines that comprise the control plane.	Array of MachinePool objects.
<code>controlPlane: architecture:</code>	Determines the instruction set architecture of the machines in the pool. Currently, clusters with varied architectures are not supported. All pools must specify the same architecture. Valid values are amd64 (the default).	String
<code>controlPlane: name:</code>	Required if you use controlPlane . The name of the machine pool.	master
<code>controlPlane: platform:</code>	Required if you use controlPlane . Use this parameter to specify the cloud provider that hosts the control plane machines. This parameter value must match the compute.platform parameter value.	aws, azure, gcp, ibmcloud, nutanix, openstack, powervs, vsphere, or {}

Parameter	Description	Values
controlPlane: replicas:	The number of control plane machines to provision.	Supported values are 3 , or 1 when deploying single-node OpenShift.
credentialsMode:	The Cloud Credential Operator (CCO) mode. If no mode is specified, the CCO dynamically tries to determine the capabilities of the provided credentials, with a preference for mint mode on the platforms where multiple modes are supported.	Mint, Passthrough, Manual or an empty string (<code>""</code>). ^[1]

Parameter	Description	Values
<p>fips:</p>	<p>Enable or disable FIPS mode. The default is false (disabled). If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.</p> <div data-bbox="486 589 592 1637" style="background-color: black; width: 66px; height: 468px; margin-bottom: 10px;"></div> <p>IMPORTANT</p> <p>To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see Installing the system in FIPS mode.</p> <p>When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.</p> <div data-bbox="486 1686 592 1883" style="background-color: black; width: 66px; height: 88px;"></div> <p>NOTE</p> <p>If you are using Azure File storage, you cannot enable FIPS mode.</p>	<p>false or true</p>


Parameter	Description	Values
<code>imageContentSources:</code>	Sources and repositories for the release-image content.	Array of objects. Includes a source and, optionally, mirrors , as described in the following rows of this table.
<code>imageContentSources: source:</code>	Required if you use imageContentSources . Specify the repository that users refer to, for example, in image pull specifications.	String
<code>imageContentSources: mirrors:</code>	Specify one or more repositories that may also contain the same images.	Array of strings
<code>publish:</code>	How to publish or expose the user-facing endpoints of your cluster, such as the Kubernetes API, OpenShift routes.	<p>Internal or External. The default value is External.</p> <p>Setting this field to Internal is not supported on non-cloud platforms.</p> <div style="display: flex; align-items: flex-start;"> <div style="width: 20px; height: 20px; background-color: black; margin-right: 5px;"></div> <div> <p>IMPORTANT</p> <p>If the value of the field is set to Internal, the cluster will become non-functional. For more information, refer to BZ#1953035.</p> </div> </div>
<code>sshKey:</code>	<p>The SSH key to authenticate access to your cluster machines.</p> <div style="display: flex; align-items: flex-start;"> <div style="width: 20px; height: 20px; background-color: black; margin-right: 5px;"></div> <div> <p>NOTE</p> <p>For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your ssh-agent process uses.</p> </div> </div>	For example, sshKey: ssh-ed25519 AAAA...

1. Not all CCO modes are supported for all cloud providers. For more information about CCO modes, see the "Managing cloud provider credentials" entry in the *Authentication and authorization* content.


23.9.1.4. Additional VMware vSphere configuration parameters

Additional VMware vSphere configuration parameters are described in the following table:

Table 23.46. Additional VMware vSphere cluster parameters

Parameter	Description	Values
platform: vsphere:	Describes your account on the cloud platform that hosts your cluster. You can use the parameter to customize the platform. If you provide additional configuration settings for compute and control plane machines in the machine pool, the parameter is not required. You can only specify one vCenter server for your OpenShift Container Platform cluster.	A dictionary of vSphere configuration objects
platform: vsphere: apiVIPs:	Virtual IP (VIP) addresses that you configured for control plane API access.  NOTE This parameter applies only to installer-provisioned infrastructure.	Multiple IP addresses
platform: vsphere: diskType:	Optional: The disk provisioning method. This value defaults to the vSphere default storage policy if not set.	Valid values are thin , thick , or eagerZeroedThick .
platform: vsphere: failureDomains:	Establishes the relationships between a region and zone. You define a failure domain by using vCenter objects, such as a datastore object. A failure domain defines the vCenter location for OpenShift Container Platform cluster nodes.	An array of failure domain configuration objects.
platform: vsphere: failureDomains: name:	The name of the failure domain.	String
platform: vsphere: failureDomains: region:	If you define multiple failure domains for your cluster, you must attach the tag to each vCenter datacenter. To define a region, use a tag from the openshift-region tag category. For a single vSphere datacenter environment, you do not need to attach a tag, but you must enter an alphanumeric value, such as datacenter , for the parameter.	String

Parameter	Description	Values
platform: vsphere: failureDomains: server:	Specifies the fully-qualified hostname or IP address of the VMware vCenter server, so that a client can access failure domain resources. You must apply the server role to the vSphere vCenter server location.	String
platform: vsphere: failureDomains: zone:	If you define multiple failure domains for your cluster, you must attach a tag to each vCenter cluster. To define a zone, use a tag from the openshift-zone tag category. For a single vSphere datacenter environment, you do not need to attach a tag, but you must enter an alphanumeric value, such as cluster , for the parameter.	String
platform: vsphere: failureDomains: topology: computeCluster:	The path to the vSphere compute cluster.	String
platform: vsphere: failureDomains: topology: datacenter:	Lists and defines the datacenters where OpenShift Container Platform virtual machines (VMs) operate. The list of datacenters must match the list of datacenters specified in the vcenters field.	String
platform: vsphere: failureDomains: topology: datastore:	Specifies the path to a vSphere datastore that stores virtual machines files for a failure domain. You must apply the datastore role to the vSphere vCenter datastore location.	String

Parameter	Description	Values
<pre>platform: vsphere: failureDomains: topology: folder:</pre>	<p>Optional: The absolute path of an existing folder where the user creates the virtual machines, for example, <code>/<datacenter_name>/vm/<folder_name>/<sub folder_name></code>. If you do not provide this value, the installation program creates a top-level folder in the datacenter virtual machine folder that is named with the infrastructure ID. If you are providing the infrastructure for the cluster and you do not want to use the default StorageClass object, named thin, you can omit the folder parameter from the install-config.yaml file.</p>	String
<pre>platform: vsphere: failureDomains: topology: networks:</pre>	<p>Lists any network in the vCenter instance that contains the virtual IP addresses and DNS records that you configured.</p>	String
<pre>platform: vsphere: failureDomains: topology: resourcePool:</pre>	<p>Optional: The absolute path of an existing resource pool where the installation program creates the virtual machines, for example, <code>/<datacenter_name>/host/<cluster_name>/Resources/<resource_pool_name>/<optional_nested_resource_pool_name></code>. If you do not specify a value, the installation program installs the resources in the root of the cluster under <code>/<datacenter_name>/host/<cluster_name>/Resources</code>.</p>	String
<pre>platform: vsphere: failureDomains: topology template:</pre>	<p>Specifies the absolute path to a pre-existing Red Hat Enterprise Linux CoreOS (RHCOS) image template or virtual machine. The installation program can use the image template or virtual machine to quickly install RHCOS on vSphere hosts. Consider using this parameter as an alternative to uploading an RHCOS image on vSphere hosts. This parameter is available for use only on installer-provisioned infrastructure.</p>	String
<pre>platform: vsphere: ingressVIPs:</pre>	<p>Virtual IP (VIP) addresses that you configured for cluster Ingress.</p> <div style="display: flex; align-items: flex-start;">  <p>NOTE</p> <p>This parameter applies only to installer-provisioned infrastructure.</p> </div>	Multiple IP addresses

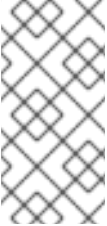

Parameter	Description	Values
platform: vsphere: vcenters:	Configures the connection details so that services can communicate with a vCenter server. Currently, only a single vCenter server is supported.	An array of vCenter configuration objects.
platform: vsphere: vcenters: datacenters:	Lists and defines the datacenters where OpenShift Container Platform virtual machines (VMs) operate. The list of datacenters must match the list of datacenters specified in the failureDomains field.	String
platform: vsphere: vcenters: password:	The password associated with the vSphere user.	String
platform: vsphere: vcenters: port:	The port number used to communicate with the vCenter server.	Integer
platform: vsphere: vcenters: server:	The fully qualified host name (FQHN) or IP address of the vCenter server.	String
platform: vsphere: vcenters: user:	The username associated with the vSphere user.	String

23.9.1.5. Deprecated VMware vSphere configuration parameters

In OpenShift Container Platform 4.13, the following vSphere configuration parameters are deprecated. You can continue to use these parameters, but the installation program does not automatically specify these parameters in the **install-config.yaml** file.

The following table lists each deprecated vSphere configuration parameter:

Table 23.47. Deprecated VMware vSphere cluster parameters

Parameter	Description	Values
platform: vsphere: apiVIP:	<p>The virtual IP (VIP) address that you configured for control plane API access.</p>  <p>NOTE</p> <p>In OpenShift Container Platform 4.12 and later, the apiVIP configuration setting is deprecated. Instead, use a List format to enter a value in the apiVIPs configuration setting.</p>	An IP address, for example 128.0.0.1 .
platform: vsphere: cluster:	The vCenter cluster to install the OpenShift Container Platform cluster in.	String
platform: vsphere: datacenter:	Defines the datacenter where OpenShift Container Platform virtual machines (VMs) operate.	String
platform: vsphere: defaultDatastore:	The name of the default datastore to use for provisioning volumes.	String
platform: vsphere: folder:	Optional: The absolute path of an existing folder where the installation program creates the virtual machines. If you do not provide this value, the installation program creates a folder that is named with the infrastructure ID in the data center virtual machine folder.	String, for example, /<datacenter_name>/vm/<folder_name>/<subfolder_name> .
platform: vsphere: ingressVIP:	<p>Virtual IP (VIP) addresses that you configured for cluster Ingress.</p>  <p>NOTE</p> <p>In OpenShift Container Platform 4.12 and later, the ingressVIP configuration setting is deprecated. Instead, use a List format to enter a value in the ingressVIPs configuration setting.</p>	An IP address, for example 128.0.0.1 .

Parameter	Description	Values
<code>platform: vsphere: network:</code>	The network in the vCenter instance that contains the virtual IP addresses and DNS records that you configured.	String
<code>platform: vsphere: password:</code>	The password for the vCenter user name.	String
<code>platform: vsphere: resourcePool:</code>	Optional: The absolute path of an existing resource pool where the installation program creates the virtual machines. If you do not specify a value, the installation program installs the resources in the root of the cluster under <code>/<datacenter_name>/host/<cluster_name>/Resources</code> .	String, for example, <code>/<datacenter_name>/host/<cluster_name>/Resources/<resource_pool_name>/<optional_nested_resource_pool_name></code> .
<code>platform: vsphere: username:</code>	The user name to use to connect to the vCenter instance with. This user must have at least the roles and privileges that are required for static or dynamic persistent volume provisioning in vSphere.	String
<code>platform: vsphere: vCenter:</code>	The fully-qualified hostname or IP address of a vCenter server.	String

23.9.1.6. Optional VMware vSphere machine pool configuration parameters

Optional VMware vSphere machine pool configuration parameters are described in the following table:

Table 23.48. Optional VMware vSphere machine pool parameters

Parameter	Description	Values
<code>platform: vsphere: clusterOSImage:</code>	The location from which the installation program downloads the Red Hat Enterprise Linux CoreOS (RHCOS) image. Before setting a path value for this parameter, ensure that the default RHCOS boot image in the OpenShift Container Platform release matches the RHCOS image template or virtual machine version; otherwise, cluster installation might fail.	An HTTP or HTTPS URL, optionally with a SHA-256 checksum. For example, <code>https://mirror.openshift.com/images/rhcos-<version>-vmware-<architecture>.ova</code> .

Parameter	Description	Values
platform: vsphere: osDisk: diskSizeGB:	The size of the disk in gigabytes.	Integer
platform: vsphere: cpus:	The total number of virtual processor cores to assign a virtual machine. The value of platform.vsphere.cpus must be a multiple of platform.vsphere.coresPerSocket value.	Integer
platform: vsphere: coresPerSocket:	The number of cores per socket in a virtual machine. The number of virtual sockets on the virtual machine is platform.vsphere.cpus/platform.vsphere.coresPerSocket . The default value for control plane nodes and worker nodes is 4 and 2 , respectively.	Integer
platform: vsphere: memoryMB:	The size of a virtual machine's memory in megabytes.	Integer

CHAPTER 24. INSTALLING ON ANY PLATFORM

24.1. INSTALLING A CLUSTER ON ANY PLATFORM

In OpenShift Container Platform version 4.16, you can install a cluster on any infrastructure that you provision, including virtualization and cloud environments.



IMPORTANT

Review the information in the [guidelines for deploying OpenShift Container Platform on non-tested platforms](#) before you attempt to install an OpenShift Container Platform cluster in virtualized or cloud environments.

24.1.1. Prerequisites

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You read the documentation on [selecting a cluster installation method and preparing it for users](#).
- If you use a firewall, you [configured it to allow the sites](#) that your cluster requires access to.



NOTE

Be sure to also review this site list if you are configuring a proxy.

24.1.2. Internet access for OpenShift Container Platform

In OpenShift Container Platform 4.16, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.



IMPORTANT

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

24.1.3. Requirements for a cluster with user-provisioned infrastructure

For a cluster that contains user-provisioned infrastructure, you must deploy all of the required machines.

This section describes the requirements for deploying OpenShift Container Platform on user-provisioned infrastructure.

24.1.3.1. Required machines for cluster installation

The smallest OpenShift Container Platform clusters require the following hosts:

Table 24.1. Minimum required hosts

Hosts	Description
One temporary bootstrap machine	The cluster requires the bootstrap machine to deploy the OpenShift Container Platform cluster on the three control plane machines. You can remove the bootstrap machine after you install the cluster.
Three control plane machines	The control plane machines run the Kubernetes and OpenShift Container Platform services that form the control plane.
At least two compute machines, which are also known as worker machines.	The workloads requested by OpenShift Container Platform users run on the compute machines.



IMPORTANT

To maintain high availability of your cluster, use separate physical hosts for these cluster machines.

The bootstrap and control plane machines must use Red Hat Enterprise Linux CoreOS (RHCOS) as the operating system. However, the compute machines can choose between Red Hat Enterprise Linux CoreOS (RHCOS), Red Hat Enterprise Linux (RHEL) 8.6 and later.

Note that RHCOS is based on Red Hat Enterprise Linux (RHEL) 9.2 and inherits all of its hardware certifications and requirements. See [Red Hat Enterprise Linux technology capabilities and limits](#).

24.1.3.2. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 24.2. Minimum resource requirements

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Bootstrap	RHCOS	4	16 GB	100 GB	300
Control plane	RHCOS	4	16 GB	100 GB	300

Machine	Operating System	vCPU [1]	Virtual RAM	Storage	Input/Output Per Second (IOPS)[2]
Compute	RHCOS, RHEL 8.6 and later [3]	2	8 GB	100 GB	300

1. One vCPU is equivalent to one physical core when simultaneous multithreading (SMT), or hyperthreading, is not enabled. When enabled, use the following formula to calculate the corresponding ratio: $(\text{threads per core} \times \text{cores}) \times \text{sockets} = \text{vCPUs}$.
2. OpenShift Container Platform and Kubernetes are sensitive to disk performance, and faster storage is recommended, particularly for etcd on the control plane nodes which require a 10 ms p99 fsync duration. Note that on many cloud platforms, storage size and IOPS scale together, so you might need to over-allocate storage volume to obtain sufficient performance.
3. As with all user-provisioned installations, if you choose to use RHEL compute machines in your cluster, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Use of RHEL 7 compute machines is deprecated and has been removed in OpenShift Container Platform 4.10 and later.



NOTE

As of OpenShift Container Platform version 4.13, RHCOS is based on RHEL version 9.2, which updates the micro-architecture requirements. The following list contains the minimum instruction set architectures (ISA) that each architecture requires:

- x86-64 architecture requires x86-64-v2 ISA
- ARM64 architecture requires ARMv8.0-A ISA
- IBM Power architecture requires Power 9 ISA
- s390x architecture requires z14 ISA

For more information, see [RHEL Architectures](#).

If an instance type for your platform meets the minimum requirements for cluster machines, it is supported to use in OpenShift Container Platform.

24.1.3.3. Certificate signing requests management

Because your cluster has limited access to automatic machine management when you use infrastructure that you provision, you must provide a mechanism for approving cluster certificate signing requests (CSRs) after installation. The **kube-controller-manager** only approves the kubelet client CSRs. The **machine-approver** cannot guarantee the validity of a serving certificate that is requested by using kubelet credentials because it cannot confirm that the correct machine issued the request. You must determine and implement a method of verifying the validity of the kubelet serving certificate requests and approving them.

24.1.3.4. Networking requirements for user-provisioned infrastructure

All the Red Hat Enterprise Linux CoreOS (RHCOS) machines require networking to be configured in **initramfs** during boot to fetch their Ignition config files.

During the initial boot, the machines require an IP address configuration that is set either through a DHCP server or statically by providing the required boot options. After a network connection is established, the machines download their Ignition config files from an HTTP or HTTPS server. The Ignition config files are then used to set the exact state of each machine. The Machine Config Operator completes more changes to the machines, such as the application of new certificates or keys, after installation.

It is recommended to use a DHCP server for long-term management of the cluster machines. Ensure that the DHCP server is configured to provide persistent IP addresses, DNS server information, and hostnames to the cluster machines.



NOTE

If a DHCP service is not available for your user-provisioned infrastructure, you can instead provide the IP networking configuration and the address of the DNS server to the nodes at RHCOS install time. These can be passed as boot arguments if you are installing from an ISO image. See the *Installing RHCOS and starting the OpenShift Container Platform bootstrap process* section for more information about static IP provisioning and advanced networking options.

The Kubernetes API server must be able to resolve the node names of the cluster machines. If the API servers and worker nodes are in different zones, you can configure a default DNS search zone to allow the API server to resolve the node names. Another supported approach is to always refer to hosts by their fully-qualified domain names in both the node objects and all DNS requests.

24.1.3.4.1. Setting the cluster node hostnames through DHCP

On Red Hat Enterprise Linux CoreOS (RHCOS) machines, the hostname is set through NetworkManager. By default, the machines obtain their hostname through DHCP. If the hostname is not provided by DHCP, set statically through kernel arguments, or another method, it is obtained through a reverse DNS lookup. Reverse DNS lookup occurs after the network has been initialized on a node and can take time to resolve. Other system services can start prior to this and detect the hostname as **localhost** or similar. You can avoid this by using DHCP to provide the hostname for each cluster node.

Additionally, setting the hostnames through DHCP can bypass any manual DNS record name configuration errors in environments that have a DNS split-horizon implementation.

24.1.3.4.2. Network connectivity requirements

You must configure the network connectivity between machines to allow OpenShift Container Platform cluster components to communicate. Each machine must be able to resolve the hostnames of all other machines in the cluster.

This section provides details about the ports that are required.



IMPORTANT

In connected OpenShift Container Platform environments, all nodes are required to have internet access to pull images for platform containers and provide telemetry data to Red Hat.

Table 24.3. Ports used for all-machine to all-machine communications

Protocol	Port	Description
ICMP	N/A	Network reachability tests
TCP	1936	Metrics
	9000-9999	Host level services, including the node exporter on ports 9100-9101 and the Cluster Version Operator on port 9099 .
	10250-10259	The default ports that Kubernetes reserves
UDP	4789	VXLAN
	6081	Geneve
	9000-9999	Host level services, including the node exporter on ports 9100-9101 .
	500	IPsec IKE packets
	4500	IPsec NAT-T packets
	123	Network Time Protocol (NTP) on UDP port 123 If an external NTP time server is configured, you must open UDP port 123 .
TCP/UDP	30000-32767	Kubernetes node port
ESP	N/A	IPsec Encapsulating Security Payload (ESP)

Table 24.4. Ports used for all-machine to control plane communications

Protocol	Port	Description
TCP	6443	Kubernetes API

Table 24.5. Ports used for control plane machine to control plane machine communications

Protocol	Port	Description
TCP	2379-2380	etcd server and peer ports

NTP configuration for user-provisioned infrastructure

OpenShift Container Platform clusters are configured to use a public Network Time Protocol (NTP) server by default. If you want to use a local enterprise NTP server, or if your cluster is being deployed in a

disconnected network, you can configure the cluster to use a specific time server. For more information, see the documentation for *Configuring chrony time service*.

If a DHCP server provides NTP server information, the chrony time service on the Red Hat Enterprise Linux CoreOS (RHCOS) machines read the information and can sync the clock with the NTP servers.

Additional resources

- [Configuring chrony time service](#)

24.1.3.5. User-provisioned DNS requirements

In OpenShift Container Platform deployments, DNS name resolution is required for the following components:

- The Kubernetes API
- The OpenShift Container Platform application wildcard
- The bootstrap, control plane, and compute machines

Reverse DNS resolution is also required for the Kubernetes API, the bootstrap machine, the control plane machines, and the compute machines.

DNS A/AAAA or CNAME records are used for name resolution and PTR records are used for reverse name resolution. The reverse records are important because Red Hat Enterprise Linux CoreOS (RHCOS) uses the reverse records to set the hostnames for all the nodes, unless the hostnames are provided by DHCP. Additionally, the reverse records are used to generate the certificate signing requests (CSR) that OpenShift Container Platform needs to operate.




NOTE

It is recommended to use a DHCP server to provide the hostnames to each cluster node. See the *DHCP recommendations for user-provisioned infrastructure* section for more information.

The following DNS records are required for a user-provisioned OpenShift Container Platform cluster and they must be in place before installation. In each record, **<cluster_name>** is the cluster name and **<base_domain>** is the base domain that you specify in the **install-config.yaml** file. A complete DNS record takes the form: **<component>.<cluster_name>.<base_domain>..**

Table 24.6. Required DNS records

Component	Record	Description
Kubernetes API	api.<cluster_name>.<base_domain>..	A DNS A/AAAA or CNAME record, and a DNS PTR record, to identify the API load balancer. These records must be resolvable by both clients external to the cluster and from all the nodes within the cluster.

Component	Record	Description
	api-int.<cluster_name>.<base_domain>	<p>A DNS A/AAAA or CNAME record, and a DNS PTR record, to internally identify the API load balancer. These records must be resolvable from all the nodes within the cluster.</p> <div style="display: flex; align-items: flex-start;">  <div> <p>IMPORTANT</p> <p>The API server must be able to resolve the worker nodes by the hostnames that are recorded in Kubernetes. If the API server cannot resolve the node names, then proxied API calls can fail, and you cannot retrieve logs from pods.</p> </div> </div>
Routes	*.apps.<cluster_name>.<base_domain>	<p>A wildcard DNS A/AAAA or CNAME record that refers to the application ingress load balancer. The application ingress load balancer targets the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default. These records must be resolvable by both clients external to the cluster and from all the nodes within the cluster.</p> <p>For example, console-openshift-console.apps.<cluster_name>.<base_domain> is used as a wildcard route to the OpenShift Container Platform console.</p>
Bootstrap machine	bootstrap.<cluster_name>.<base_domain>	A DNS A/AAAA or CNAME record, and a DNS PTR record, to identify the bootstrap machine. These records must be resolvable by the nodes within the cluster.
Control plane machines	<control_plane><n>.<cluster_name>.<base_domain>	DNS A/AAAA or CNAME records and DNS PTR records to identify each machine for the control plane nodes. These records must be resolvable by the nodes within the cluster.
Compute machines	<compute><n>.<cluster_name>.<base_domain>	DNS A/AAAA or CNAME records and DNS PTR records to identify each machine for the worker nodes. These records must be resolvable by the nodes within the cluster.

**NOTE**

In OpenShift Container Platform 4.4 and later, you do not need to specify etcd host and SRV records in your DNS configuration.

TIP

You can use the **dig** command to verify name and reverse name resolution. See the section on *Validating DNS resolution for user-provisioned infrastructure* for detailed validation steps.

24.1.3.5.1. Example DNS configuration for user-provisioned clusters

This section provides A and PTR record configuration samples that meet the DNS requirements for deploying OpenShift Container Platform on user-provisioned infrastructure. The samples are not meant to provide advice for choosing one DNS solution over another.

In the examples, the cluster name is **ocp4** and the base domain is **example.com**.

Example DNS A record configuration for a user-provisioned cluster

The following example is a BIND zone file that shows sample A records for name resolution in a user-provisioned cluster.

Example 24.1. Sample DNS zone database

```
$TTL 1W
@ IN SOA ns1.example.com. root (
  2019070700 ; serial
  3H ; refresh (3 hours)
  30M ; retry (30 minutes)
  2W ; expiry (2 weeks)
  1W ) ; minimum (1 week)
IN NS ns1.example.com.
IN MX 10 smtp.example.com.
;
;
ns1.example.com. IN A 192.168.1.5
smtp.example.com. IN A 192.168.1.5
;
helper.example.com. IN A 192.168.1.5
helper.ocp4.example.com. IN A 192.168.1.5
;
api.ocp4.example.com. IN A 192.168.1.5 1
api-int.ocp4.example.com. IN A 192.168.1.5 2
;
*.apps.ocp4.example.com. IN A 192.168.1.5 3
;
bootstrap.ocp4.example.com. IN A 192.168.1.96 4
;
control-plane0.ocp4.example.com. IN A 192.168.1.97 5
control-plane1.ocp4.example.com. IN A 192.168.1.98 6
control-plane2.ocp4.example.com. IN A 192.168.1.99 7
;
compute0.ocp4.example.com. IN A 192.168.1.11 8
compute1.ocp4.example.com. IN A 192.168.1.7 9
;
;EOF
```

- 1 Provides name resolution for the Kubernetes API. The record refers to the IP address of the API load balancer.
- 2 Provides name resolution for the Kubernetes API. The record refers to the IP address of the API load balancer and is used for internal cluster communications.
- 3 Provides name resolution for the wildcard routes. The record refers to the IP address of the application ingress load balancer. The application ingress load balancer targets the machines

application ingress load balancer. The application ingress load balancer targets the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.



NOTE

In the example, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

- 4 Provides name resolution for the bootstrap machine.
- 5 6 7 Provides name resolution for the control plane machines.
- 8 9 Provides name resolution for the compute machines.

Example DNS PTR record configuration for a user-provisioned cluster

The following example BIND zone file shows sample PTR records for reverse name resolution in a user-provisioned cluster.

Example 24.2. Sample DNS zone database for reverse records

```
$TTL 1W
@ IN SOA ns1.example.com. root (
    2019070700 ; serial
    3H ; refresh (3 hours)
    30M ; retry (30 minutes)
    2W ; expiry (2 weeks)
    1W ) ; minimum (1 week)
IN NS ns1.example.com.
;
5.1.168.192.in-addr.arpa. IN PTR api.ocp4.example.com. 1
5.1.168.192.in-addr.arpa. IN PTR api-int.ocp4.example.com. 2
;
96.1.168.192.in-addr.arpa. IN PTR bootstrap.ocp4.example.com. 3
;
97.1.168.192.in-addr.arpa. IN PTR control-plane0.ocp4.example.com. 4
98.1.168.192.in-addr.arpa. IN PTR control-plane1.ocp4.example.com. 5
99.1.168.192.in-addr.arpa. IN PTR control-plane2.ocp4.example.com. 6
;
11.1.168.192.in-addr.arpa. IN PTR compute0.ocp4.example.com. 7
7.1.168.192.in-addr.arpa. IN PTR compute1.ocp4.example.com. 8
;
;EOF
```

- 1 Provides reverse DNS resolution for the Kubernetes API. The PTR record refers to the record name of the API load balancer.
- 2 Provides reverse DNS resolution for the Kubernetes API. The PTR record refers to the record name of the API load balancer and is used for internal cluster communications.

3 Provides reverse DNS resolution for the bootstrap machine.

4 5 6 Provides reverse DNS resolution for the control plane machines.

7 8 Provides reverse DNS resolution for the compute machines.



NOTE

A PTR record is not required for the OpenShift Container Platform application wildcard.

24.1.3.6. Load balancing requirements for user-provisioned infrastructure

Before you install OpenShift Container Platform, you must provision the API and application Ingress load balancing infrastructure. In production scenarios, you can deploy the API and application Ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.



NOTE

If you want to deploy the API and application Ingress load balancers with a Red Hat Enterprise Linux (RHEL) instance, you must purchase the RHEL subscription separately.

The load balancing infrastructure must meet the following requirements:

1. **API load balancer.** Provides a common endpoint for users, both human and machine, to interact with and configure the platform. Configure the following conditions:
 - Layer 4 load balancing only. This can be referred to as Raw TCP or SSL Passthrough mode.
 - A stateless load balancing algorithm. The options vary based on the load balancer implementation.



IMPORTANT

Do not configure session persistence for an API load balancer. Configuring session persistence for a Kubernetes API server might cause performance issues from excess application traffic for your OpenShift Container Platform cluster and the Kubernetes API that runs inside the cluster.

Configure the following ports on both the front and back of the load balancers:

Table 24.7. API load balancer

Port	Back-end machines (pool members)	Internal	External	Description
6443	Bootstrap and control plane. You remove the bootstrap machine from the load balancer after the bootstrap machine initializes the cluster control plane. You must configure the /readyz endpoint for the API server health check probe.	X	X	Kubernetes API server

Port	Back-end machines (pool members)	Internal	External	Description
22623	Bootstrap and control plane. You remove the bootstrap machine from the load balancer after the bootstrap machine initializes the cluster control plane.	X		Machine config server



NOTE

The load balancer must be configured to take a maximum of 30 seconds from the time the API server turns off the `/readyz` endpoint to the removal of the API server instance from the pool. Within the time frame after `/readyz` returns an error or becomes healthy, the endpoint must have been removed or added. Probing every 5 or 10 seconds, with two successful requests to become healthy and three to become unhealthy, are well-tested values.

2. **Application Ingress load balancer.** Provides an ingress point for application traffic flowing in from outside the cluster. A working configuration for the Ingress router is required for an OpenShift Container Platform cluster.

Configure the following conditions:

- Layer 4 load balancing only. This can be referred to as Raw TCP or SSL Passthrough mode.
- A connection-based or session-based persistence is recommended, based on the options available and types of applications that will be hosted on the platform.

TIP

If the true IP address of the client can be seen by the application Ingress load balancer, enabling source IP-based session persistence can improve performance for applications that use end-to-end TLS encryption.

Configure the following ports on both the front and back of the load balancers:

Table 24.8. Application Ingress load balancer

Port	Back-end machines (pool members)	Internal	External	Description
443	The machines that run the Ingress Controller pods, compute, or worker, by default.	X	X	HTTPS traffic
80	The machines that run the Ingress Controller pods, compute, or worker, by default.	X	X	HTTP traffic

**NOTE**

If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application Ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes.

24.1.3.6.1. Example load balancer configuration for user-provisioned clusters

This section provides an example API and application Ingress load balancer configuration that meets the load balancing requirements for user-provisioned clusters. The sample is an `/etc/haproxy/haproxy.cfg` configuration for an HAProxy load balancer. The example is not meant to provide advice for choosing one load balancing solution over another.

In the example, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

**NOTE**

If you are using HAProxy as a load balancer and SELinux is set to **enforcing**, you must ensure that the HAProxy service can bind to the configured TCP port by running **setsebool -P haproxy_connect_any=1**.

Example 24.3. Sample API and application Ingress load balancer configuration

```

global
  log      127.0.0.1 local2
  pidfile  /var/run/haproxy.pid
  maxconn  4000
  daemon
defaults
  mode            http
  log             global
  option          dontlognull
  option http-server-close
  option          redispatch
  retries         3
  timeout http-request  10s
  timeout queue       1m
  timeout connect    10s
  timeout client     1m
  timeout server     1m
  timeout http-keep-alive 10s
  timeout check      10s
  maxconn           3000
listen api-server-6443 1
  bind *:6443
  mode tcp
  option httpchk GET /readyz HTTP/1.0
  option log-health-checks
  balance roundrobin
  server bootstrap bootstrap.ocp4.example.com:6443 verify none check check-ssl inter 10s fall 2
  rise 3 backup 2
  server master0 master0.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s


```

```

fall 2 rise 3
server master1 master1.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
fall 2 rise 3
server master2 master2.ocp4.example.com:6443 weight 1 verify none check check-ssl inter 10s
fall 2 rise 3
listen machine-config-server-22623 3
bind *:22623
mode tcp
server bootstrap bootstrap.ocp4.example.com:22623 check inter 1s backup 4
server master0 master0.ocp4.example.com:22623 check inter 1s
server master1 master1.ocp4.example.com:22623 check inter 1s
server master2 master2.ocp4.example.com:22623 check inter 1s
listen ingress-router-443 5
bind *:443
mode tcp
balance source
server compute0 compute0.ocp4.example.com:443 check inter 1s
server compute1 compute1.ocp4.example.com:443 check inter 1s
listen ingress-router-80 6
bind *:80
mode tcp
balance source
server compute0 compute0.ocp4.example.com:80 check inter 1s
server compute1 compute1.ocp4.example.com:80 check inter 1s

```

- 1 Port **6443** handles the Kubernetes API traffic and points to the control plane machines.
- 2 4 The bootstrap entries must be in place before the OpenShift Container Platform cluster installation and they must be removed after the bootstrap process is complete.
- 3 Port **22623** handles the machine config server traffic and points to the control plane machines.
- 5 Port **443** handles the HTTPS traffic and points to the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.
- 6 Port **80** handles the HTTP traffic and points to the machines that run the Ingress Controller pods. The Ingress Controller pods run on the compute machines by default.



NOTE

If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application Ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes.

TIP

If you are using HAProxy as a load balancer, you can check that the **haproxy** process is listening on ports **6443**, **22623**, **443**, and **80** by running **netstat -nltp** on the HAProxy node.

24.1.4. Preparing the user-provisioned infrastructure

Before you install OpenShift Container Platform on user-provisioned infrastructure, you must prepare the underlying infrastructure.

This section provides details about the high-level steps required to set up your cluster infrastructure in preparation for an OpenShift Container Platform installation. This includes configuring IP networking and network connectivity for your cluster nodes, enabling the required ports through your firewall, and setting up the required DNS and load balancing infrastructure.

After preparation, your cluster infrastructure must meet the requirements outlined in the *Requirements for a cluster with user-provisioned infrastructure* section.

Prerequisites

- You have reviewed the [OpenShift Container Platform 4.x Tested Integrations](#) page.
- You have reviewed the infrastructure requirements detailed in the *Requirements for a cluster with user-provisioned infrastructure* section.

Procedure

1. If you are using DHCP to provide the IP networking configuration to your cluster nodes, configure your DHCP service.
 - a. Add persistent IP addresses for the nodes to your DHCP server configuration. In your configuration, match the MAC address of the relevant network interface to the intended IP address for each node.
 - b. When you use DHCP to configure IP addressing for the cluster machines, the machines also obtain the DNS server information through DHCP. Define the persistent DNS server address that is used by the cluster nodes through your DHCP server configuration.



NOTE

If you are not using a DHCP service, you must provide the IP networking configuration and the address of the DNS server to the nodes at RHCOS install time. These can be passed as boot arguments if you are installing from an ISO image. See the *Installing RHCOS and starting the OpenShift Container Platform bootstrap process* section for more information about static IP provisioning and advanced networking options.

- c. Define the hostnames of your cluster nodes in your DHCP server configuration. See the *Setting the cluster node hostnames through DHCP* section for details about hostname considerations.



NOTE

If you are not using a DHCP service, the cluster nodes obtain their hostname through a reverse DNS lookup.

2. Ensure that your network infrastructure provides the required network connectivity between the cluster components. See the *Networking requirements for user-provisioned infrastructure* section for details about the requirements.

3. Configure your firewall to enable the ports required for the OpenShift Container Platform cluster components to communicate. See *Networking requirements for user-provisioned infrastructure* section for details about the ports that are required.



IMPORTANT

By default, port **1936** is accessible for an OpenShift Container Platform cluster, because each control plane node needs access to this port.

Avoid using the Ingress load balancer to expose this port, because doing so might result in the exposure of sensitive information, such as statistics and metrics, related to Ingress Controllers.

4. Setup the required DNS infrastructure for your cluster.
 - a. Configure DNS name resolution for the Kubernetes API, the application wildcard, the bootstrap machine, the control plane machines, and the compute machines.
 - b. Configure reverse DNS resolution for the Kubernetes API, the bootstrap machine, the control plane machines, and the compute machines.
See the *User-provisioned DNS requirements* section for more information about the OpenShift Container Platform DNS requirements.
5. Validate your DNS configuration.
 - a. From your installation node, run DNS lookups against the record names of the Kubernetes API, the wildcard routes, and the cluster nodes. Validate that the IP addresses in the responses correspond to the correct components.
 - b. From your installation node, run reverse DNS lookups against the IP addresses of the load balancer and the cluster nodes. Validate that the record names in the responses correspond to the correct components.
See the *Validating DNS resolution for user-provisioned infrastructure* section for detailed DNS validation steps.
6. Provision the required API and application ingress load balancing infrastructure. See the *Load balancing requirements for user-provisioned infrastructure* section for more information about the requirements.



NOTE

Some load balancing solutions require the DNS name resolution for the cluster nodes to be in place before the load balancing is initialized.

24.1.5. Validating DNS resolution for user-provisioned infrastructure

You can validate your DNS configuration before installing OpenShift Container Platform on user-provisioned infrastructure.



IMPORTANT

The validation steps detailed in this section must succeed before you install your cluster.

Prerequisites

- You have configured the required DNS records for your user-provisioned infrastructure.

Procedure

- From your installation node, run DNS lookups against the record names of the Kubernetes API, the wildcard routes, and the cluster nodes. Validate that the IP addresses contained in the responses correspond to the correct components.
 - Perform a lookup against the Kubernetes API record name. Check that the result points to the IP address of the API load balancer:

```
$ dig +noall +answer @<nameserver_ip> api.<cluster_name>.<base_domain> 1
```

- Replace **<nameserver_ip>** with the IP address of the nameserver, **<cluster_name>** with your cluster name, and **<base_domain>** with your base domain name.

Example output

```
api.ocp4.example.com. 604800 IN A 192.168.1.5
```

- Perform a lookup against the Kubernetes internal API record name. Check that the result points to the IP address of the API load balancer:

```
$ dig +noall +answer @<nameserver_ip> api-int.<cluster_name>.<base_domain>
```

Example output

```
api-int.ocp4.example.com. 604800 IN A 192.168.1.5
```

- Test an example ***.apps.<cluster_name>.<base_domain>** DNS wildcard lookup. All of the application wildcard lookups must resolve to the IP address of the application ingress load balancer:

```
$ dig +noall +answer @<nameserver_ip> random.apps.<cluster_name>.<base_domain>
```

Example output

```
random.apps.ocp4.example.com. 604800 IN A 192.168.1.5
```



NOTE

In the example outputs, the same load balancer is used for the Kubernetes API and application ingress traffic. In production scenarios, you can deploy the API and application ingress load balancers separately so that you can scale the load balancer infrastructure for each in isolation.

You can replace **random** with another wildcard value. For example, you can query the route to the OpenShift Container Platform console:

```
$ dig +noall +answer @<nameserver_ip> console-openshift-console.apps.<cluster_name>.<base_domain>
```

Example output

```
console-openshift-console.apps.ocp4.example.com. 604800 IN A 192.168.1.5
```

- d. Run a lookup against the bootstrap DNS record name. Check that the result points to the IP address of the bootstrap node:

```
$ dig +noall +answer @<nameserver_ip> bootstrap.<cluster_name>.<base_domain>
```

Example output

```
bootstrap.ocp4.example.com. 604800 IN A 192.168.1.96
```

- e. Use this method to perform lookups against the DNS record names for the control plane and compute nodes. Check that the results correspond to the IP addresses of each node.
2. From your installation node, run reverse DNS lookups against the IP addresses of the load balancer and the cluster nodes. Validate that the record names contained in the responses correspond to the correct components.
 - a. Perform a reverse lookup against the IP address of the API load balancer. Check that the response includes the record names for the Kubernetes API and the Kubernetes internal API:

```
$ dig +noall +answer @<nameserver_ip> -x 192.168.1.5
```

Example output

```
5.1.168.192.in-addr.arpa. 604800 IN PTR api-int.ocp4.example.com. 1  
5.1.168.192.in-addr.arpa. 604800 IN PTR api.ocp4.example.com. 2
```

1 Provides the record name for the Kubernetes internal API.

2 Provides the record name for the Kubernetes API.



NOTE

A PTR record is not required for the OpenShift Container Platform application wildcard. No validation step is needed for reverse DNS resolution against the IP address of the application ingress load balancer.

- b. Perform a reverse lookup against the IP address of the bootstrap node. Check that the result points to the DNS record name of the bootstrap node:

```
$ dig +noall +answer @<nameserver_ip> -x 192.168.1.96
```

Example output

```
96.1.168.192.in-addr.arpa. 604800 IN PTR bootstrap.ocp4.example.com.
```

- c. Use this method to perform reverse lookups against the IP addresses for the control plane and compute nodes. Check that the results correspond to the DNS record names of each node.

24.1.6. Generating a key pair for cluster node SSH access

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.



IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.



NOTE

You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.



NOTE

If you plan to install an OpenShift Container Platform cluster that uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.



NOTE

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the `ssh-agent` process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

Example output

```
Agent pid 31874
```



NOTE

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the `ssh-agent`:

```
$ ssh-add <path>/<file_name> 1
```

- 1** Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program. If you install a cluster on infrastructure that you provision, you must provide the key to the installation program.

24.1.7. Obtaining the installation program


Before you install OpenShift Container Platform, download the installation file on the host you are using for installation.

Prerequisites

- You have a computer that runs Linux or macOS, with 500 MB of local disk space.


Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select your infrastructure provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.



IMPORTANT

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.



IMPORTANT

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.


4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from Red Hat OpenShift Cluster Manager](#). This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.

24.1.8. Installing the OpenShift CLI

You can install the OpenShift CLI (**oc**) to interact with OpenShift Container Platform from a command-line interface. You can install **oc** on Linux, Windows, or macOS.



IMPORTANT

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.16. Download and install the new version of **oc**.

Installing the OpenShift CLI on Linux

You can install the OpenShift CLI (**oc**) binary on Linux by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.16 Linux Client** entry and save the file.
5. Unpack the archive:

```
$ tar xvf <file>
```

6. Place the **oc** binary in a directory that is on your **PATH**.
To check your **PATH**, execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

Installing the OpenShift CLI on Windows

You can install the OpenShift CLI (**oc**) binary on Windows by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.
To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
C:\> oc <command>
```

Installing the OpenShift CLI on macOS

You can install the OpenShift CLI (**oc**) binary on macOS by using the following procedure.

Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.16 macOS Client** entry and save the file.

**NOTE**

For macOS arm64, choose the **OpenShift v4.16 macOS arm64 Client** entry.

4. Unpack and unzip the archive.
5. Move the **oc** binary to a directory on your PATH.
To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification

- After you install the OpenShift CLI, it is available using the **oc** command:

```
$ oc <command>
```

24.1.9. Manually creating the installation configuration file

Installing the cluster requires that you manually create the installation configuration file.

Prerequisites

- You have an SSH public key on your local machine to provide to the installation program. The key will be used for SSH authentication onto your cluster nodes for debugging and disaster recovery.
- You have obtained the OpenShift Container Platform installation program and the pull secret for your cluster.

Procedure


1. Create an installation directory to store your required installation assets in:

```
$ mkdir <installation_directory>
```

**IMPORTANT**

You must create a directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

2. Customize the sample **install-config.yaml** file template that is provided and save it in the **<installation_directory>**.

**NOTE**

You must name this configuration file **install-config.yaml**.

3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.

**IMPORTANT**

The **install-config.yaml** file is consumed during the next step of the installation process. You must back it up now.

24.1.9.1. Sample install-config.yaml file for other platforms

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.

```

apiVersion: v1
baseDomain: example.com 1
compute: 2
- hyperthreading: Enabled 3
  name: worker
  replicas: 0 4
controlPlane: 5
  hyperthreading: Enabled 6
  name: master
  replicas: 3 7
metadata:
  name: test 8
networking:
  clusterNetwork:
  - cidr: 10.128.0.0/14 9
    hostPrefix: 23 10
  networkType: OVNKubernetes 11
  serviceNetwork: 12
  - 172.30.0.0/16
platform:
  none: {} 13
fips: false 14
pullSecret: '{"auths": ...}' 15
sshKey: 'ssh-ed25519 AAAA...' 16

```

- 1 The base domain of the cluster. All DNS records must be sub-domains of this base and include the cluster name.
- 2 5 The **controlPlane** section is a single mapping, but the **compute** section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Only one control plane pool is used.

- 3 6** Specifies whether to enable or disable simultaneous multithreading (SMT), or hyperthreading. By default, SMT is enabled to increase the performance of the cores in your machines. You can



NOTE

Simultaneous multithreading (SMT) is enabled by default. If SMT is not enabled in your BIOS settings, the **hyperthreading** parameter has no effect.



IMPORTANT

If you disable **hyperthreading**, whether in the BIOS or in the **install-config.yaml** file, ensure that your capacity planning accounts for the dramatically decreased machine performance.

- 4** You must set this value to **0** when you install OpenShift Container Platform on user-provisioned infrastructure. In installer-provisioned installations, the parameter controls the number of compute machines that the cluster creates and manages for you. In user-provisioned installations, you must manually deploy the compute machines before you finish installing the cluster.



NOTE

If you are installing a three-node cluster, do not deploy any compute machines when you install the Red Hat Enterprise Linux CoreOS (RHCOS) machines.

- 7** The number of control plane machines that you add to the cluster. Because the cluster uses these values as the number of etcd endpoints in the cluster, the value must match the number of control plane machines that you deploy.
- 8** The cluster name that you specified in your DNS records.
- 9** A block of IP addresses from which pod IP addresses are allocated. This block must not overlap with existing physical networks. These IP addresses are used for the pod network. If you need to access the pods from an external network, you must configure load balancers and routers to manage the traffic.



NOTE

Class E CIDR range is reserved for a future use. To use the Class E CIDR range, you must ensure your networking environment accepts the IP addresses within the Class E CIDR range.

- 10** The subnet prefix length to assign to each individual node. For example, if **hostPrefix** is set to **23**, then each node is assigned a **/23** subnet out of the given **cidr**, which allows for 510 ($2^{(32 - 23)} - 2$) pod IP addresses. If you are required to provide access to nodes from an external network, configure load balancers and routers to manage the traffic.
- 11** The cluster network plugin to install. The default value **OVNKubernetes** is the only supported value.
- 12** The IP address pool to use for service IP addresses. You can enter only one IP address pool. This block must not overlap with existing physical networks. If you need to access the services from an external network, configure load balancers and routers to manage the traffic.
- 13** You must set the platform to **none**. You cannot provide additional platform configuration variables

for your platform.



IMPORTANT

Clusters that are installed with the platform type **none** are unable to use some features, such as managing compute machines with the Machine API. This limitation applies even if the compute machines that are attached to the cluster are installed on a platform that would normally support the feature. This parameter cannot be changed after installation.

- 14 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.



IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

- 15 The [pull secret from Red Hat OpenShift Cluster Manager](#) . This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.
- 16 The SSH public key for the **core** user in Red Hat Enterprise Linux CoreOS (RHCOS).



NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.


24.1.9.2. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to

bypass the proxy if necessary.



NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use ***** to bypass the proxy for all destinations.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.

**NOTE**

The installation program does not support the proxy **readinessEndpoints** field.

**NOTE**

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

```
$. /openshift-install wait-for install-complete --log-level debug
```

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.

**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

24.1.9.3. Configuring a three-node cluster

Optionally, you can deploy zero compute machines in a bare metal cluster that consists of three control plane machines only. This provides smaller, more resource efficient clusters for cluster administrators and developers to use for testing, development, and production.

In three-node OpenShift Container Platform environments, the three control plane machines are schedulable, which means that your application workloads are scheduled to run on them.

Prerequisites

- You have an existing **install-config.yaml** file.

Procedure

- Ensure that the number of compute replicas is set to **0** in your **install-config.yaml** file, as shown in the following **compute** stanza:

```
compute:
- name: worker
  platform: {}
  replicas: 0
```



NOTE

You must set the value of the **replicas** parameter for the compute machines to **0** when you install OpenShift Container Platform on user-provisioned infrastructure, regardless of the number of compute machines you are deploying. In installer-provisioned installations, the parameter controls the number of compute machines that the cluster creates and manages for you. This does not apply to user-provisioned installations, where the compute machines are deployed manually.


For three-node cluster installations, follow these next steps:

- If you are deploying a three-node cluster with zero compute nodes, the Ingress Controller pods run on the control plane nodes. In three-node cluster deployments, you must configure your application ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes. See the *Load balancing requirements for user-provisioned infrastructure* section for more information.
- When you create the Kubernetes manifest files in the following procedure, ensure that the **mastersSchedulable** parameter in the `<installation_directory>/manifests/cluster-scheduler-02-config.yml` file is set to **true**. This enables your application workloads to run on the control plane nodes.
- Do not deploy any compute nodes when you create the Red Hat Enterprise Linux CoreOS (RHCOS) machines.

24.1.10. Creating the Kubernetes manifest and Ignition config files

Because you must modify some cluster definition files and manually start the cluster machines, you must generate the Kubernetes manifest and Ignition config files that the cluster needs to configure the machines.

The installation configuration file transforms into the Kubernetes manifests. The manifests wrap into the Ignition configuration files, which are later used to configure the cluster machines.



IMPORTANT

- The Ignition config files that the OpenShift Container Platform installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

Prerequisites

- You obtained the OpenShift Container Platform installation program.

- You created the **install-config.yaml** installation configuration file.

Procedure

- Change to the directory that contains the OpenShift Container Platform installation program and generate the Kubernetes manifests for the cluster:

```
$ ./openshift-install create manifests --dir <installation_directory> 1
```

- For **<installation_directory>**, specify the installation directory that contains the **install-config.yaml** file you created.



WARNING

If you are installing a three-node cluster, skip the following step to allow the control plane nodes to be schedulable.



IMPORTANT

When you configure control plane nodes from the default unschedulable to schedulable, additional subscriptions are required. This is because control plane nodes then become compute nodes.


- Check that the **mastersSchedulable** parameter in the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** Kubernetes manifest file is set to **false**. This setting prevents pods from being scheduled on the control plane machines:
 - Open the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** file.
 - Locate the **mastersSchedulable** parameter and ensure that it is set to **false**.
 - Save and exit the file.
- To create the Ignition configuration files, run the following command from the directory that contains the installation program:

```
$ ./openshift-install create ignition-configs --dir <installation_directory> 1
```

- For **<installation_directory>**, specify the same installation directory.

Ignition config files are created for the bootstrap, control plane, and compute nodes in the installation directory. The **kubeadmin-password** and **kubeconfig** files are created in the **./<installation_directory>/auth** directory:

```
.
├── auth
│   ├── kubeadmin-password
```

24.1.11. Installing RHCOS and starting the OpenShift Container Platform bootstrap process

To install OpenShift Container Platform on bare metal infrastructure that you provision, you must install Red Hat Enterprise Linux CoreOS (RHCOS) on the machines. When you install RHCOS, you must provide the Ignition config file that was generated by the OpenShift Container Platform installation program for the type of machine you are installing. If you have configured suitable networking, DNS, and load balancing infrastructure, the OpenShift Container Platform bootstrap process begins automatically after the RHCOS machines have rebooted.

To install RHCOS on the machines, follow either the steps to use an ISO image or network PXE booting.



NOTE

The compute node deployment steps included in this installation document are RHCOS-specific. If you choose instead to deploy RHEL-based compute nodes, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks. Only RHEL 8 compute machines are supported.

You can configure RHCOS during ISO and PXE installations by using the following methods:

- **Kernel arguments:** You can use kernel arguments to provide installation-specific information. For example, you can specify the locations of the RHCOS installation files that you uploaded to your HTTP server and the location of the Ignition config file for the type of node you are installing. For a PXE installation, you can use the **APPEND** parameter to pass the arguments to the kernel of the live installer. For an ISO installation, you can interrupt the live installation boot process to add the kernel arguments. In both installation cases, you can use special **coreos.inst.*** arguments to direct the live installer, as well as standard installation boot arguments for turning standard kernel services on or off.
- **Ignition configs:** OpenShift Container Platform Ignition config files (***.ign**) are specific to the type of node you are installing. You pass the location of a bootstrap, control plane, or compute node Ignition config file during the RHCOS installation so that it takes effect on first boot. In special cases, you can create a separate, limited Ignition config to pass to the live system. That Ignition config could do a certain set of tasks, such as reporting success to a provisioning system after completing installation. This special Ignition config is consumed by the **coreos-installer** to be applied on first boot of the installed system. Do not provide the standard control plane and compute node Ignition configs to the live ISO directly.
- **coreos-installer:** You can boot the live ISO installer to a shell prompt, which allows you to prepare the permanent system in a variety of ways before first boot. In particular, you can run the **coreos-installer** command to identify various artifacts to include, work with disk partitions, and set up networking. In some cases, you can configure features on the live system and copy them to the installed system.

Whether to use an ISO or PXE install depends on your situation. A PXE install requires an available DHCP service and more preparation, but can make the installation process more automated. An ISO install is a more manual process and can be inconvenient if you are setting up more than a few machines.

**NOTE**

As of OpenShift Container Platform 4.6, the RHCOS ISO and other installation artifacts provide support for installation on disks with 4K sectors.

24.1.11.1. Installing RHCOS by using an ISO image

You can use an ISO image to install RHCOS on the machines.

Prerequisites

- You have created the Ignition config files for your cluster.
- You have configured suitable network, DNS and load balancing infrastructure.
- You have an HTTP server that can be accessed from your computer, and from the machines that you create.
- You have reviewed the *Advanced RHCOS installation configuration* section for different ways to configure features, such as networking and disk partitioning.


Procedure

1. Obtain the SHA512 digest for each of your Ignition config files. For example, you can use the following on a system running Linux to get the SHA512 digest for your **bootstrap.ign** Ignition config file:

```
$ sha512sum <installation_directory>/bootstrap.ign
```

The digests are provided to the **coreos-installer** in a later step to validate the authenticity of the Ignition config files on the cluster nodes.

2. Upload the bootstrap, control plane, and compute node Ignition config files that the installation program created to your HTTP server. Note the URLs of these files.

**IMPORTANT**

You can add or change configuration settings in your Ignition configs before saving them to your HTTP server. If you plan to add more compute machines to your cluster after you finish installation, do not delete these files.

3. From the installation host, validate that the Ignition config files are available on the URLs. The following example gets the Ignition config file for the bootstrap node:

```
$ curl -k http://<HTTP_server>/bootstrap.ign 1
```

Example output

```
% Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
           Dload  Upload  Total   Spent    Left    Speed
  0   0   0   0   0   0   0   0  0 --:--:-- --:--:-- --:--:--    0{"ignition":
{"version":"3.2.0"},"passwd":{"users":[{"name":"core","sshAuthorizedKeys":["ssh-rsa...
```

Replace **bootstrap.ign** with **master.ign** or **worker.ign** in the command to validate that the Ignition config files for the control plane and compute nodes are also available.

- Although it is possible to obtain the RHCOS images that are required for your preferred method of installing operating system instances from the [RHCOS image mirror](#) page, the recommended way to obtain the correct version of your RHCOS images are from the output of **openshift-install** command:

```
$ openshift-install coreos print-stream-json | grep '\.iso[^\.]'
```

Example output

```
"location": "<url>/art/storage/releases/rhcos-4.16-aarch64/<release>/aarch64/rhcos-
<release>-live.aarch64.iso",
"location": "<url>/art/storage/releases/rhcos-4.16-ppc64le/<release>/ppc64le/rhcos-
<release>-live.ppc64le.iso",
"location": "<url>/art/storage/releases/rhcos-4.16-s390x/<release>/s390x/rhcos-<release>-
live.s390x.iso",
"location": "<url>/art/storage/releases/rhcos-4.16/<release>/x86_64/rhcos-<release>-
live.x86_64.iso",
```



IMPORTANT

The RHCOS images might not change with every release of OpenShift Container Platform. You must download images with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Use the image versions that match your OpenShift Container Platform version if they are available. Use only ISO images for this procedure. RHCOS qcow2 images are not supported for this installation type.

ISO file names resemble the following example:

rhcos-<version>-live.<architecture>.iso

- Use the ISO to start the RHCOS installation. Use one of the following installation options:
 - Burn the ISO image to a disk and boot it directly.
 - Use ISO redirection by using a lights-out management (LOM) interface.
- Boot the RHCOS ISO image without specifying any options or interrupting the live boot sequence. Wait for the installer to boot into a shell prompt in the RHCOS live environment.



NOTE

It is possible to interrupt the RHCOS installation boot process to add kernel arguments. However, for this ISO procedure you should use the **coreos-installer** command as outlined in the following steps, instead of adding kernel arguments.

- Run the **coreos-installer** command and specify the options that meet your installation requirements. At a minimum, you must specify the URL that points to the Ignition config file for the node type, and the device that you are installing to:

```
$ sudo coreos-installer install --ignition-url=http://<HTTP_server>/<node_type>.ign <device>
--ignition-hash=sha512-<digest> 1 2
```

1 1 You must run the **coreos-installer** command by using **sudo**, because the **core** user does not have the required root privileges to perform the installation.

2 The **--ignition-hash** option is required when the Ignition config file is obtained through an HTTP URL to validate the authenticity of the Ignition config file on the cluster node. **<digest>** is the Ignition config file SHA512 digest obtained in a preceding step.




NOTE

If you want to provide your Ignition config files through an HTTPS server that uses TLS, you can add the internal certificate authority (CA) to the system trust store before running **coreos-installer**.

The following example initializes a bootstrap node installation to the **/dev/sda** device. The Ignition config file for the bootstrap node is obtained from an HTTP web server with the IP address 192.168.1.2:

```
$ sudo coreos-installer install --ignition-
url=http://192.168.1.2:80/installation_directory/bootstrap.ign /dev/sda --ignition-hash=sha512-
a5a2d43879223273c9b60af66b44202a1d1248fc01cf156c46d4a79f552b6bad47bc8cc78ddf011
6e80c59d2ea9e32ba53bc807afbca581aa059311def2c3e3b
```

8. Monitor the progress of the RHCOS installation on the console of the machine.



IMPORTANT

Be sure that the installation is successful on each node before commencing with the OpenShift Container Platform installation. Observing the installation process can also help to determine the cause of RHCOS installation issues that might arise.


9. After RHCOS installs, you must reboot the system. During the system reboot, it applies the Ignition config file that you specified.

10. Check the console output to verify that Ignition ran.

Example command

```
Ignition: ran on 2022/03/14 14:48:33 UTC (this boot)
Ignition: user-provided config was applied
```

11. Continue to create the other machines for your cluster.



IMPORTANT

You must create the bootstrap and control plane machines at this time. If the control plane machines are not made schedulable, also create at least two compute machines before you install OpenShift Container Platform.

If the required network, DNS, and load balancer infrastructure are in place, the OpenShift Container Platform bootstrap process begins automatically after the RHCOS nodes have rebooted.



NOTE

RHCOS nodes do not include a default password for the **core** user. You can access the nodes by running **ssh core@<node>.<cluster_name>.<base_domain>** as a user with access to the SSH private key that is paired to the public key that you specified in your **install_config.yaml** file. OpenShift Container Platform 4 cluster nodes running RHCOS are immutable and rely on Operators to apply cluster changes. Accessing cluster nodes by using SSH is not recommended. However, when investigating installation issues, if the OpenShift Container Platform API is not available, or the kubelet is not properly functioning on a target node, SSH access might be required for debugging or disaster recovery.

24.1.11.2. Installing RHCOS by using PXE or iPXE booting

You can use PXE or iPXE booting to install RHCOS on the machines.

Prerequisites

- You have created the Ignition config files for your cluster.
- You have configured suitable network, DNS and load balancing infrastructure.
- You have configured suitable PXE or iPXE infrastructure.
- You have an HTTP server that can be accessed from your computer, and from the machines that you create.
- You have reviewed the *Advanced RHCOS installation configuration* section for different ways to configure features, such as networking and disk partitioning.

Procedure

1. Upload the bootstrap, control plane, and compute node Ignition config files that the installation program created to your HTTP server. Note the URLs of these files.



IMPORTANT

You can add or change configuration settings in your Ignition configs before saving them to your HTTP server. If you plan to add more compute machines to your cluster after you finish installation, do not delete these files.

2. From the installation host, validate that the Ignition config files are available on the URLs. The following example gets the Ignition config file for the bootstrap node:

```
$ curl -k http://<HTTP_server>/bootstrap.ign 1
```

Example output

```
% Total % Received % Xferd Average Speed Time Time Time Current
```

```

Dload Upload Total Spent Left Speed
0 0 0 0 0 0 0 0 --:--:-- --:--:-- --:--:-- 0{"ignition":
{"version":"3.2.0"},"passwd":{"users":[{"name":"core","sshAuthorizedKeys":["ssh-rsa...

```

Replace **bootstrap.ign** with **master.ign** or **worker.ign** in the command to validate that the Ignition config files for the control plane and compute nodes are also available.

- Although it is possible to obtain the RHCOS **kernel**, **initramfs** and **rootfs** files that are required for your preferred method of installing operating system instances from the [RHCOS image mirror](#) page, the recommended way to obtain the correct version of your RHCOS files are from the output of **openshift-install** command:

```
$ openshift-install coreos print-stream-json | grep -Eo "https.*(kernel-|initramfs-|rootfs-)|w+(\.img)?"
```

Example output

```

"<url>/art/storage/releases/rhcos-4.16-aarch64/<release>/aarch64/rhcos-<release>-live-
kernel-aarch64"
"<url>/art/storage/releases/rhcos-4.16-aarch64/<release>/aarch64/rhcos-<release>-live-
initramfs.aarch64.img"
"<url>/art/storage/releases/rhcos-4.16-aarch64/<release>/aarch64/rhcos-<release>-live-
rootfs.aarch64.img"
"<url>/art/storage/releases/rhcos-4.16-ppc64le/49.84.202110081256-0/ppc64le/rhcos-
<release>-live-kernel-ppc64le"
"<url>/art/storage/releases/rhcos-4.16-ppc64le/<release>/ppc64le/rhcos-<release>-live-
initramfs.ppc64le.img"
"<url>/art/storage/releases/rhcos-4.16-ppc64le/<release>/ppc64le/rhcos-<release>-live-
rootfs.ppc64le.img"
"<url>/art/storage/releases/rhcos-4.16-s390x/<release>/s390x/rhcos-<release>-live-kernel-
s390x"
"<url>/art/storage/releases/rhcos-4.16-s390x/<release>/s390x/rhcos-<release>-live-
initramfs.s390x.img"
"<url>/art/storage/releases/rhcos-4.16-s390x/<release>/s390x/rhcos-<release>-live-
rootfs.s390x.img"
"<url>/art/storage/releases/rhcos-4.16/<release>/x86_64/rhcos-<release>-live-kernel-
x86_64"
"<url>/art/storage/releases/rhcos-4.16/<release>/x86_64/rhcos-<release>-live-
initramfs.x86_64.img"
"<url>/art/storage/releases/rhcos-4.16/<release>/x86_64/rhcos-<release>-live-
rootfs.x86_64.img"

```



IMPORTANT

The RHCOS artifacts might not change with every release of OpenShift Container Platform. You must download images with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Only use the appropriate **kernel**, **initramfs**, and **rootfs** artifacts described below for this procedure. RHCOS QCOW2 images are not supported for this installation type.

The file names contain the OpenShift Container Platform version number. They resemble the following examples:

- **kernel:** `rhcos-<version>-live-kernel-<architecture>`
- **initramfs:** `rhcos-<version>-live-initramfs.<architecture>.img`
- **rootfs:** `rhcos-<version>-live-rootfs.<architecture>.img`

4. Upload the **rootfs**, **kernel**, and **initramfs** files to your HTTP server.



IMPORTANT

If you plan to add more compute machines to your cluster after you finish installation, do not delete these files.

5. Configure the network boot infrastructure so that the machines boot from their local disks after RHCOS is installed on them.
6. Configure PXE or iPXE installation for the RHCOS images and begin the installation. Modify one of the following example menu entries for your environment and verify that the image and Ignition files are properly accessible:

- For PXE (**x86_64**):

```

DEFAULT pxeboot
TIMEOUT 20
PROMPT 0
LABEL pxeboot
  KERNEL http://<HTTP_server>/rhcos-<version>-live-kernel-<architecture> 1
  APPEND initrd=http://<HTTP_server>/rhcos-<version>-live-initramfs.
<architecture>.img coreos.live.rootfs_url=http://<HTTP_server>/rhcos-<version>-live-
rootfs.<architecture>.img coreos.inst.install_dev=/dev/sda
coreos.inst.ignition_url=http://<HTTP_server>/bootstrap.ign 2 3

```

- 1 1** Specify the location of the live **kernel** file that you uploaded to your HTTP server. The URL must be HTTP, TFTP, or FTP; HTTPS and NFS are not supported.
- 2** If you use multiple NICs, specify a single interface in the **ip** option. For example, to use DHCP on a NIC that is named **eno1**, set **ip=eno1:dhcp**.
- 3** Specify the locations of the RHCOS files that you uploaded to your HTTP server. The **initrd** parameter value is the location of the **initramfs** file, the **coreos.live.rootfs_url** parameter value is the location of the **rootfs** file, and the **coreos.inst.ignition_url** parameter value is the location of the bootstrap Ignition config file. You can also add more kernel arguments to the **APPEND** line to configure networking or other boot options.

**NOTE**

This configuration does not enable serial console access on machines with a graphical console. To configure a different console, add one or more **console=** arguments to the **APPEND** line. For example, add **console=tty0 console=ttyS0** to set the first PC serial port as the primary console and the graphical console as a secondary console. For more information, see [How does one set up a serial terminal and/or console in Red Hat Enterprise Linux?](#) and "Enabling the serial console for PXE and ISO installation" in the "Advanced RHCOS installation configuration" section.

- For iPXE (**x86_64** + **aarch64**):

```
kernel http://<HTTP_server>/rhcos-<version>-live-kernel-<architecture> initrd=main
coreos.live.rootfs_url=http://<HTTP_server>/rhcos-<version>-live-rootfs.
<architecture>.img coreos.inst.install_dev=/dev/sda
coreos.inst.ignition_url=http://<HTTP_server>/bootstrap.ign 1 2
initrd --name main http://<HTTP_server>/rhcos-<version>-live-initramfs.
<architecture>.img 3
boot
```

- 1 Specify the locations of the RHCOS files that you uploaded to your HTTP server. The **kernel** parameter value is the location of the **kernel** file, the **initrd=main** argument is needed for booting on UEFI systems, the **coreos.live.rootfs_url** parameter value is the location of the **rootfs** file, and the **coreos.inst.ignition_url** parameter value is the location of the bootstrap Ignition config file.
- 2 If you use multiple NICs, specify a single interface in the **ip** option. For example, to use DHCP on a NIC that is named **eno1**, set **ip=eno1:dhcp**.
- 3 Specify the location of the **initramfs** file that you uploaded to your HTTP server.

**NOTE**

This configuration does not enable serial console access on machines with a graphical console. To configure a different console, add one or more **console=** arguments to the **kernel** line. For example, add **console=tty0 console=ttyS0** to set the first PC serial port as the primary console and the graphical console as a secondary console. For more information, see [How does one set up a serial terminal and/or console in Red Hat Enterprise Linux?](#) and "Enabling the serial console for PXE and ISO installation" in the "Advanced RHCOS installation configuration" section.

**NOTE**

To network boot the CoreOS **kernel** on **aarch64** architecture, you need to use a version of iPXE build with the **IMAGE_GZIP** option enabled. See [IMAGE_GZIP option in iPXE](#).

- For PXE (with UEFI and Grub as second stage) on **aarch64**:

```
menuentry 'Install CoreOS' {
    linux rhcos-<version>-live-kernel-<architecture>
```




```

coreos.live.rootfs_url=http://<HTTP_server>/rhcos-<version>-live-rootfs.
<architecture>.img coreos.inst.install_dev=/dev/sda
coreos.inst.ignition_url=http://<HTTP_server>/bootstrap.ign 1 2
  initrd rhcos-<version>-live-initramfs.<architecture>.img 3
}

```

- 1** Specify the locations of the RHCOS files that you uploaded to your HTTP/TFTP server. The **kernel** parameter value is the location of the **kernel** file on your TFTP server. The **coreos.live.rootfs_url** parameter value is the location of the **rootfs** file, and the **coreos.inst.ignition_url** parameter value is the location of the bootstrap Ignition config file on your HTTP Server.
 - 2** If you use multiple NICs, specify a single interface in the **ip** option. For example, to use DHCP on a NIC that is named **eno1**, set **ip=eno1:dhcp**.
 - 3** Specify the location of the **initramfs** file that you uploaded to your TFTP server.
7. Monitor the progress of the RHCOS installation on the console of the machine.



IMPORTANT

Be sure that the installation is successful on each node before commencing with the OpenShift Container Platform installation. Observing the installation process can also help to determine the cause of RHCOS installation issues that might arise.

8. After RHCOS installs, the system reboots. During reboot, the system applies the Ignition config file that you specified.
9. Check the console output to verify that Ignition ran.


Example command

```

Ignition: ran on 2022/03/14 14:48:33 UTC (this boot)
Ignition: user-provided config was applied

```

10. Continue to create the machines for your cluster.



IMPORTANT

You must create the bootstrap and control plane machines at this time. If the control plane machines are not made schedulable, also create at least two compute machines before you install the cluster.

If the required network, DNS, and load balancer infrastructure are in place, the OpenShift Container Platform bootstrap process begins automatically after the RHCOS nodes have rebooted.



NOTE

RHCOS nodes do not include a default password for the **core** user. You can access the nodes by running **ssh core@<node>.<cluster_name>.<base_domain>** as a user with access to the SSH private key that is paired to the public key that you specified in your **install_config.yaml** file. OpenShift Container Platform 4 cluster nodes running RHCOS are immutable and rely on Operators to apply cluster changes. Accessing cluster nodes by using SSH is not recommended. However, when investigating installation issues, if the OpenShift Container Platform API is not available, or the kubelet is not properly functioning on a target node, SSH access might be required for debugging or disaster recovery.

24.1.11.3. Advanced RHCOS installation configuration

A key benefit for manually provisioning the Red Hat Enterprise Linux CoreOS (RHCOS) nodes for OpenShift Container Platform is to be able to do configuration that is not available through default OpenShift Container Platform installation methods. This section describes some of the configurations that you can do using techniques that include:

- Passing kernel arguments to the live installer
- Running **coreos-installer** manually from the live system
- Customizing a live ISO or PXE boot image

The advanced configuration topics for manual Red Hat Enterprise Linux CoreOS (RHCOS) installations detailed in this section relate to disk partitioning, networking, and using Ignition configs in different ways.

24.1.11.3.1. Using advanced networking options for PXE and ISO installations

Networking for OpenShift Container Platform nodes uses DHCP by default to gather all necessary configuration settings. To set up static IP addresses or configure special settings, such as bonding, you can do one of the following:

- Pass special kernel parameters when you boot the live installer.
- Use a machine config to copy networking files to the installed system.
- Configure networking from a live installer shell prompt, then copy those settings to the installed system so that they take effect when the installed system first boots.

To configure a PXE or iPXE installation, use one of the following options:

- See the "Advanced RHCOS installation reference" tables.
- Use a machine config to copy networking files to the installed system.

To configure an ISO installation, use the following procedure.

Procedure

1. Boot the ISO installer.
2. From the live system shell prompt, configure networking for the live system using available RHEL tools, such as **nmcli** or **nmtui**.

- Run the **coreos-installer** command to install the system, adding the **--copy-network** option to copy networking configuration. For example:

```
$ sudo coreos-installer install --copy-network \
  --ignition-url=http://host/worker.ign /dev/disk/by-id/scsi-<serial_number>
```



IMPORTANT

The **--copy-network** option only copies networking configuration found under **/etc/NetworkManager/system-connections**. In particular, it does not copy the system hostname.


- Reboot into the installed system.

Additional resources

- See [Getting started with nmcli](#) and [Getting started with nmtui](#) in the RHEL 8 documentation for more information about the **nmcli** and **nmtui** tools.

24.1.11.3.2. Disk partitioning

Disk partitions are created on OpenShift Container Platform cluster nodes during the Red Hat Enterprise Linux CoreOS (RHCOS) installation. Each RHCOS node of a particular architecture uses the same partition layout, unless you override the default partitioning configuration. During the RHCOS installation, the size of the root file system is increased to use any remaining available space on the target device.



IMPORTANT

The use of a custom partition scheme on your node might result in OpenShift Container Platform not monitoring or alerting on some node partitions. If you override the default partitioning, see [Understanding OpenShift File System Monitoring \(eviction conditions\)](#) for more information about how OpenShift Container Platform monitors your host file systems.

OpenShift Container Platform monitors the following two filesystem identifiers:

- nodefs**, which is the filesystem that contains **/var/lib/kubelet**
- imagefs**, which is the filesystem that contains **/var/lib/containers**

For the default partition scheme, **nodefs** and **imagefs** monitor the same root filesystem, **/**.

To override the default partitioning when installing RHCOS on an OpenShift Container Platform cluster node, you must create separate partitions. Consider a situation where you want to add a separate storage partition for your containers and container images. For example, by mounting **/var/lib/containers** in a separate partition, the kubelet separately monitors **/var/lib/containers** as the **imagefs** directory and the root file system as the **nodefs** directory.



IMPORTANT


If you have resized your disk size to host a larger file system, consider creating a separate **/var/lib/containers** partition. Consider resizing a disk that has an **xfs** format to reduce CPU time issues caused by a high number of allocation groups.

24.1.11.3.2.1. Creating a separate `/var` partition

In general, you should use the default disk partitioning that is created during the RHCOS installation. However, there are cases where you might want to create a separate partition for a directory that you expect to grow.

OpenShift Container Platform supports the addition of a single partition to attach storage to either the `/var` directory or a subdirectory of `/var`. For example:

- `/var/lib/containers`: Holds container-related content that can grow as more images and containers are added to a system.
- `/var/lib/etcd`: Holds data that you might want to keep separate for purposes such as performance optimization of etcd storage.
- `/var`: Holds data that you might want to keep separate for purposes such as auditing.



IMPORTANT

For disk sizes larger than 100GB, and especially larger than 1TB, create a separate `/var` partition.

Storing the contents of a `/var` directory separately makes it easier to grow storage for those areas as needed and reinstall OpenShift Container Platform at a later date and keep that data intact. With this method, you will not have to pull all your containers again, nor will you have to copy massive log files when you update systems.

The use of a separate partition for the `/var` directory or a subdirectory of `/var` also prevents data growth in the partitioned directory from filling up the root file system.

The following procedure sets up a separate `/var` partition by adding a machine config manifest that is wrapped into the Ignition config file for a node type during the preparation phase of an installation.

Procedure

1. On your installation host, change to the directory that contains the OpenShift Container Platform installation program and generate the Kubernetes manifests for the cluster:

```
$ openshift-install create manifests --dir <installation_directory>
```

2. Create a Butane config that configures the additional partition. For example, name the file `$HOME/clusterconfig/98-var-partition.bu`, change the disk device name to the name of the storage device on the `worker` systems, and set the storage size as appropriate. This example places the `/var` directory on a separate partition:

```
variant: openshift
version: 4.16.0
metadata:
  labels:
    machineconfiguration.openshift.io/role: worker
  name: 98-var-partition
storage:
  disks:
    - device: /dev/disk/by-id/<device_name> 1
    partitions:
```

```

- label: var
  start_mib: <partition_start_offset> 2
  size_mib: <partition_size> 3
  number: 5
filesystems:
- device: /dev/disk/by-partlabel/var
  path: /var
  format: xfs
  mount_options: [defaults, prjquota] 4
  with_mount_unit: true

```

- 1 The storage device name of the disk that you want to partition.
- 2 When adding a data partition to the boot disk, a minimum offset value of 25000 mebibytes is recommended. The root file system is automatically resized to fill all available space up to the specified offset. If no offset value is specified, or if the specified value is smaller than the recommended minimum, the resulting root file system will be too small, and future reinstalls of RHCOS might overwrite the beginning of the data partition.
- 3 The size of the data partition in mebibytes.
- 4 The **prjquota** mount option must be enabled for filesystems used for container storage.



NOTE

When creating a separate **/var** partition, you cannot use different instance types for compute nodes, if the different instance types do not have the same device name.

3. Create a manifest from the Butane config and save it to the **clusterconfig/openshift** directory. For example, run the following command:

```
$ butane $HOME/clusterconfig/98-var-partition.bu -o $HOME/clusterconfig/openshift/98-var-partition.yaml
```

4. Create the Ignition config files:

```
$ openshift-install create ignition-configs --dir <installation_directory> 1
```

- 1 For **<installation_directory>**, specify the same installation directory.

Ignition config files are created for the bootstrap, control plane, and compute nodes in the installation directory:

```

.
├── auth
│   ├── kubeadmin-password
│   └── kubeconfig
├── bootstrap.ign
├── master.ign
├── metadata.json
└── worker.ign

```

The files in the `<installation_directory>/manifest` and `<installation_directory>/openshift` directories are wrapped into the Ignition config files, including the file that contains the **98-var-partition** custom **MachineConfig** object.

Next steps

- You can apply the custom disk partitioning by referencing the Ignition config files during the RHCOS installations.

24.1.11.3.2.2. Retaining existing partitions

For an ISO installation, you can add options to the **coreos-installer** command that cause the installer to maintain one or more existing partitions. For a PXE installation, you can add **coreos.inst.*** options to the **APPEND** parameter to preserve partitions.

Saved partitions might be data partitions from an existing OpenShift Container Platform system. You can identify the disk partitions you want to keep either by partition label or by number.



NOTE

If you save existing partitions, and those partitions do not leave enough space for RHCOS, the installation will fail without damaging the saved partitions.

Retaining existing partitions during an ISO installation

This example preserves any partition in which the partition label begins with **data** (**data***):

```
# coreos-installer install --ignition-url http://10.0.2.2:8080/user.ign \
  --save-partlabel 'data*' /dev/disk/by-id/scsi-<serial_number>
```

The following example illustrates running the **coreos-installer** in a way that preserves the sixth (6) partition on the disk:

```
# coreos-installer install --ignition-url http://10.0.2.2:8080/user.ign \
  --save-partindex 6 /dev/disk/by-id/scsi-<serial_number>
```

This example preserves partitions 5 and higher:

```
# coreos-installer install --ignition-url http://10.0.2.2:8080/user.ign
  --save-partindex 5- /dev/disk/by-id/scsi-<serial_number>
```

In the previous examples where partition saving is used, **coreos-installer** recreates the partition immediately.

Retaining existing partitions during a PXE installation

This **APPEND** option preserves any partition in which the partition label begins with 'data' ('data*'):

```
coreos.inst.save_partlabel=data*
```

This **APPEND** option preserves partitions 5 and higher:

```
coreos.inst.save_partindex=5-
```


This **APPEND** option preserves partition 6:

```
coreos.inst.save_partindex=6
```

24.1.11.3.3. Identifying Ignition configs

When doing an RHCOS manual installation, there are two types of Ignition configs that you can provide, with different reasons for providing each one:

- **Permanent install Ignition config:** Every manual RHCOS installation needs to pass one of the Ignition config files generated by **openshift-installer**, such as **bootstrap.ign**, **master.ign** and **worker.ign**, to carry out the installation.



IMPORTANT

It is not recommended to modify these Ignition config files directly. You can update the manifest files that are wrapped into the Ignition config files, as outlined in examples in the preceding sections.

For PXE installations, you pass the Ignition configs on the **APPEND** line using the **coreos.inst.ignition_url=** option. For ISO installations, after the ISO boots to the shell prompt, you identify the Ignition config on the **coreos-installer** command line with the **--ignition-url=** option. In both cases, only HTTP and HTTPS protocols are supported.

- **Live install Ignition config:** This type can be created by using the **coreos-installer customize** subcommand and its various options. With this method, the Ignition config passes to the live install medium, runs immediately upon booting, and performs setup tasks before or after the RHCOS system installs to disk. This method should only be used for performing tasks that must be done once and not applied again later, such as with advanced partitioning that cannot be done using a machine config.


For PXE or ISO boots, you can create the Ignition config and **APPEND** the **ignition.config.url=** option to identify the location of the Ignition config. You also need to append **ignition.firstboot** or **ignition.platform.id=metal** or the **ignition.config.url** option will be ignored.

24.1.11.3.4. Advanced RHCOS installation reference

This section illustrates the networking configuration and other advanced options that allow you to modify the Red Hat Enterprise Linux CoreOS (RHCOS) manual installation process. The following tables describe the kernel arguments and command-line options you can use with the RHCOS live installer and the **coreos-installer** command.

24.1.11.3.4.1. Networking and bonding options for ISO installations

If you install RHCOS from an ISO image, you can add kernel arguments manually when you boot the image to configure networking for a node. If no networking arguments are specified, DHCP is activated in the initramfs when RHCOS detects that networking is required to fetch the Ignition config file.



IMPORTANT

When adding networking arguments manually, you must also add the **rd.neednet=1** kernel argument to bring the network up in the initramfs.

The following information provides examples for configuring networking and bonding on your RHCOS nodes for ISO installations. The examples describe how to use the **ip=**, **nameserver=**, and **bond=** kernel arguments.



NOTE

Ordering is important when adding the kernel arguments: **ip=**, **nameserver=**, and then **bond=**.

The networking options are passed to the **dracut** tool during system boot. For more information about the networking options supported by **dracut**, see the [dracut.cmdline manual page](#).

The following examples are the networking options for ISO installation.

Configuring DHCP or static IP addresses

To configure an IP address, either use DHCP (**ip=dhcp**) or set an individual static IP address (**ip=<host_ip>**). If setting a static IP, you must then identify the DNS server IP address (**nameserver=<dns_ip>**) on each node. The following example sets:

- The node's IP address to **10.10.10.2**
- The gateway address to **10.10.10.254**
- The netmask to **255.255.255.0**
- The hostname to **core0.example.com**
- The DNS server address to **4.4.4.41**
- The auto-configuration value to **none**. No auto-configuration is required when IP networking is configured statically.

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none  
nameserver=4.4.4.41
```



NOTE

When you use DHCP to configure IP addressing for the RHCOS machines, the machines also obtain the DNS server information through DHCP. For DHCP-based deployments, you can define the DNS server address that is used by the RHCOS nodes through your DHCP server configuration.

Configuring an IP address without a static hostname

You can configure an IP address without assigning a static hostname. If a static hostname is not set by the user, it will be picked up and automatically set by a reverse DNS lookup. To configure an IP address without a static hostname refer to the following example:

- The node's IP address to **10.10.10.2**
- The gateway address to **10.10.10.254**
- The netmask to **255.255.255.0**
- The DNS server address to **4.4.4.41**

- The auto-configuration value to **none**. No auto-configuration is required when IP networking is configured statically.

```
ip=10.10.10.2::10.10.10.254:255.255.255.0::enp1s0:none
nameserver=4.4.4.41
```

Specifying multiple network interfaces

You can specify multiple network interfaces by setting multiple **ip=** entries.

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none
ip=10.10.10.3::10.10.10.254:255.255.255.0:core0.example.com:enp2s0:none
```

Configuring default gateway and route

Optional: You can configure routes to additional networks by setting an **rd.route=** value.



NOTE

When you configure one or multiple networks, one default gateway is required. If the additional network gateway is different from the primary network gateway, the default gateway must be the primary network gateway.

- Run the following command to configure the default gateway:

```
ip=::10.10.10.254:::
```

- Enter the following command to configure the route for the additional network:

```
rd.route=20.20.20.0/24:20.20.20.254:enp2s0
```

Disabling DHCP on a single interface

You can disable DHCP on a single interface, such as when there are two or more network interfaces and only one interface is being used. In the example, the **enp1s0** interface has a static networking configuration and DHCP is disabled for **enp2s0**, which is not used:

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp1s0:none
ip=::::core0.example.com:enp2s0:none
```

Combining DHCP and static IP configurations

You can combine DHCP and static IP configurations on systems with multiple network interfaces, for example:

```
ip=enp1s0:dhcp
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp2s0:none
```

Configuring VLANs on individual interfaces

Optional: You can configure VLANs on individual interfaces by using the **vlan=** parameter.

- To configure a VLAN on a network interface and use a static IP address, run the following command:

```
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:enp2s0.100:none
vlan=enp2s0.100:enp2s0
```

- To configure a VLAN on a network interface and to use DHCP, run the following command:

```
ip=enp2s0.100:dhcp
vlan=enp2s0.100:enp2s0
```

Providing multiple DNS servers

You can provide multiple DNS servers by adding a **nameserver=** entry for each server, for example:

```
nameserver=1.1.1.1
nameserver=8.8.8.8
```

Bonding multiple network interfaces to a single interface

Optional: You can bond multiple network interfaces to a single interface by using the **bond=** option. Refer to the following examples:


- The syntax for configuring a bonded interface is: **bond=<name>[:<network_interfaces>] [:<options>]**
<name> is the bonding device name (**bond0**), **<network_interfaces>** represents a comma-separated list of physical (ethernet) interfaces (**em1,em2**), and *options* is a comma-separated list of bonding options. Enter **modinfo bonding** to see available options.
- When you create a bonded interface using **bond=**, you must specify how the IP address is assigned and other information for the bonded interface.
 - To configure the bonded interface to use DHCP, set the bond's IP address to **dhcp**. For example:

```
bond=bond0:em1,em2:mode=active-backup
ip=bond0:dhcp
```

- To configure the bonded interface to use a static IP address, enter the specific IP address you want and related information. For example:

```
bond=bond0:em1,em2:mode=active-backup
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:bond0:none
```

Bonding multiple SR-IOV network interfaces to a dual port NIC interface



IMPORTANT

Support for Day 1 operations associated with enabling NIC partitioning for SR-IOV devices is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

Optional: You can bond multiple SR-IOV network interfaces to a dual port NIC interface by using the **bond=** option.

On each node, you must perform the following tasks:

1. Create the SR-IOV virtual functions (VFs) following the guidance in [Managing SR-IOV devices](#). Follow the procedure in the "Attaching SR-IOV networking devices to virtual machines" section.
2. Create the bond, attach the desired VFs to the bond and set the bond link state up following the guidance in [Configuring network bonding](#). Follow any of the described procedures to create the bond.

The following examples illustrate the syntax you must use:

- The syntax for configuring a bonded interface is **bond=<name>[:<network_interfaces>][:options]**. **<name>** is the bonding device name (**bond0**), **<network_interfaces>** represents the virtual functions (VFs) by their known name in the kernel and shown in the output of the **ip link** command (**eno1f0**, **eno2f0**), and **options** is a comma-separated list of bonding options. Enter **modinfo bonding** to see available options.
- When you create a bonded interface using **bond=**, you must specify how the IP address is assigned and other information for the bonded interface.
 - To configure the bonded interface to use DHCP, set the bond's IP address to **dhcp**. For example:

```
bond=bond0:eno1f0,eno2f0:mode=active-backup
ip=bond0:dhcp
```

- To configure the bonded interface to use a static IP address, enter the specific IP address you want and related information. For example:

```
bond=bond0:eno1f0,eno2f0:mode=active-backup
ip=10.10.10.2::10.10.10.254:255.255.255.0:core0.example.com:bond0:none
```

Using network teaming

Optional: You can use a network teaming as an alternative to bonding by using the **team=** parameter:

- The syntax for configuring a team interface is: **team=name[:network_interfaces]**. **name** is the team device name (**team0**) and **network_interfaces** represents a comma-separated list of physical (ethernet) interfaces (**em1**, **em2**).



NOTE

Teaming is planned to be deprecated when RHCOS switches to an upcoming version of RHEL. For more information, see this [Red Hat Knowledgebase Article](#).

Use the following example to configure a network team:

```
team=team0:em1,em2
ip=team0:dhcp
```


24.11.3.4.2. coreos-installer options for ISO and PXE installations

You can install RHCOS by running **coreos-installer install <options> <device>** at the command prompt, after booting into the RHCOS live environment from an ISO image.

The following table shows the subcommands, options, and arguments you can pass to the **coreos-installer** command.


Table 24.9. **coreos-installer** subcommands, command-line options, and arguments

coreos-installer install subcommand	
<i>Subcommand</i>	<i>Description</i>
\$ coreos-installer install <options> <device>	Embed an Ignition config in an ISO image.
coreos-installer install subcommand options	
<i>Option</i>	<i>Description</i>
-u, --image-url <url>	Specify the image URL manually.
-f, --image-file <path>	Specify a local image file manually. Used for debugging.
-i, --ignition-file <path>	Embed an Ignition config from a file.
-l, --ignition-url <URL>	Embed an Ignition config from a URL.
--ignition-hash <digest>	Digest type-value of the Ignition config.
-p, --platform <name>	Override the Ignition platform ID for the installed system.
--console <spec>	Set the kernel and bootloader console for the installed system. For more information about the format of <spec> , see the Linux kernel serial console documentation.
--append-karg <arg>...	Append a default kernel argument to the installed system.
--delete-karg <arg>...	Delete a default kernel argument from the installed system.

-n, --copy-network	Copy the network configuration from the install environment.
	 <p>IMPORTANT</p> <p>The --copy-network option only copies networking configuration found under /etc/NetworkManager/system-connections. In particular, it does not copy the system hostname.</p>
--network-dir <path>	For use with -n . Default is /etc/NetworkManager/system-connections/ .
--save-partlabel <lx>..	Save partitions with this label glob.
--save-partindex <id>...	Save partitions with this number or range.
--insecure	Skip RHCOS image signature verification.
--insecure-ignition	Allow Ignition URL without HTTPS or hash.
--architecture <name>	Target CPU architecture. Valid values are x86_64 and aarch64 .
--preserve-on-error	Do not clear partition table on error.
-h, --help	Print help information.
coreos-installer install subcommand argument	
<i>Argument</i>	<i>Description</i>
<device>	The destination device.
coreos-installer ISO subcommands	
<i>Subcommand</i>	<i>Description</i>
\$ coreos-installer iso customize <options> <ISO_image>	Customize a RHCOS live ISO image.
coreos-installer iso reset <options> <ISO_image>	Restore a RHCOS live ISO image to default settings.

coreos-installer iso ignition remove <options> <ISO_image>	Remove the embedded Ignition config from an ISO image.
coreos-installer ISO customize subcommand options	
<i>Option</i>	<i>Description</i>
--dest-ignition <path>	Merge the specified Ignition config file into a new configuration fragment for the destination system.
--dest-console <spec>	Specify the kernel and bootloader console for the destination system.
--dest-device <path>	Install and overwrite the specified destination device.
--dest-karg-append <arg>	Add a kernel argument to each boot of the destination system.
--dest-karg-delete <arg>	Delete a kernel argument from each boot of the destination system.
--network-keyfile <path>	Configure networking by using the specified NetworkManager keyfile for live and destination systems.
--ignition-ca <path>	Specify an additional TLS certificate authority to be trusted by Ignition.
--pre-install <path>	Run the specified script before installation.
--post-install <path>	Run the specified script after installation.
--installer-config <path>	Apply the specified installer configuration file.
--live-ignition <path>	Merge the specified Ignition config file into a new configuration fragment for the live environment.
--live-karg-append <arg>	Add a kernel argument to each boot of the live environment.
--live-karg-delete <arg>	Delete a kernel argument from each boot of the live environment.
--live-karg-replace <k=o=n>	Replace a kernel argument in each boot of the live environment, in the form key=old=new .
-f, --force	Overwrite an existing Ignition config.

-o, --output <path>	Write the ISO to a new output file.
-h, --help	Print help information.
coreos-installer PXE subcommands	
<i>Subcommand</i>	<i>Description</i>
Note that not all of these options are accepted by all subcommands.	
coreos-installer pxe customize <options> <path>	Customize a RHCOS live PXE boot config.
coreos-installer pxe ignition wrap <options>	Wrap an Ignition config in an image.
coreos-installer pxe ignition unwrap <options> <image_name>	Show the wrapped Ignition config in an image.
coreos-installer PXE customize subcommand options	
<i>Option</i>	<i>Description</i>
Note that not all of these options are accepted by all subcommands.	
--dest-ignition <path>	Merge the specified Ignition config file into a new configuration fragment for the destination system.
--dest-console <spec>	Specify the kernel and bootloader console for the destination system.
--dest-device <path>	Install and overwrite the specified destination device.
--network-keyfile <path>	Configure networking by using the specified NetworkManager keyfile for live and destination systems.
--ignition-ca <path>	Specify an additional TLS certificate authority to be trusted by Ignition.
--pre-install <path>	Run the specified script before installation.
post-install <path>	Run the specified script after installation.
--installer-config <path>	Apply the specified installer configuration file.
--live-ignition <path>	Merge the specified Ignition config file into a new configuration fragment for the live environment.

-o, --output <path>	Write the initramfs to a new output file.  NOTE This option is required for PXE environments.
-h, --help	Print help information.

24.1.11.3.4.3. **coreos.inst** boot options for ISO or PXE installations

You can automatically invoke **coreos-installer** options at boot time by passing **coreos.inst** boot arguments to the RHCOS live installer. These are provided in addition to the standard boot arguments.

- For ISO installations, the **coreos.inst** options can be added by interrupting the automatic boot at the bootloader menu. You can interrupt the automatic boot by pressing **TAB** while the **RHEL CoreOS (Live)** menu option is highlighted.
- For PXE or iPXE installations, the **coreos.inst** options must be added to the **APPEND** line before the RHCOS live installer is booted.

The following table shows the RHCOS live installer **coreos.inst** boot options for ISO and PXE installations.

Table 24.10. **coreos.inst** boot options

Argument	Description
coreos.inst.install_dev	Required. The block device on the system to install to. It is recommended to use the full path, such as /dev/sda , although sda is allowed.
coreos.inst.ignition_url	Optional: The URL of the Ignition config to embed into the installed system. If no URL is specified, no Ignition config is embedded. Only HTTP and HTTPS protocols are supported.
coreos.inst.save_partlabel	Optional: Comma-separated labels of partitions to preserve during the install. Glob-style wildcards are permitted. The specified partitions do not need to exist.
coreos.inst.save_partindex	Optional: Comma-separated indexes of partitions to preserve during the install. Ranges m-n are permitted, and either m or n can be omitted. The specified partitions do not need to exist.
coreos.inst.insecure	Optional: Permits the OS image that is specified by coreos.inst.image_url to be unsigned.

Argument	Description
coreos.inst.image_url	<p>Optional: Download and install the specified RHCOS image.</p> <ul style="list-style-type: none"> • This argument should not be used in production environments and is intended for debugging purposes only. • While this argument can be used to install a version of RHCOS that does not match the live media, it is recommended that you instead use the media that matches the version you want to install. • If you are using coreos.inst.image_url, you must also use coreos.inst.insecure. This is because the bare-metal media are not GPG-signed for OpenShift Container Platform. • Only HTTP and HTTPS protocols are supported.
coreos.inst.skip_reboot	<p>Optional: The system will not reboot after installing. After the install finishes, you will receive a prompt that allows you to inspect what is happening during installation. This argument should not be used in production environments and is intended for debugging purposes only.</p>
coreos.inst.platform_id	<p>Optional: The Ignition platform ID of the platform the RHCOS image is being installed on. Default is metal. This option determines whether or not to request an Ignition config from the cloud provider, such as VMware. For example: coreos.inst.platform_id=vmware.</p>
ignition.config.url	<p>Optional: The URL of the Ignition config for the live boot. For example, this can be used to customize how coreos-installer is invoked, or to run code before or after the installation. This is different from coreos.inst.ignition_url, which is the Ignition config for the installed system.</p>

24.1.12. Waiting for the bootstrap process to complete

The OpenShift Container Platform bootstrap process begins after the cluster nodes first boot into the persistent RHCOS environment that has been installed to disk. The configuration information provided through the Ignition config files is used to initialize the bootstrap process and install OpenShift Container Platform on the machines. You must wait for the bootstrap process to complete.

Prerequisites

- You have created the Ignition config files for your cluster.
- You have configured suitable network, DNS and load balancing infrastructure.
- You have obtained the installation program and generated the Ignition config files for your cluster.
- You installed RHCOS on your cluster machines and provided the Ignition config files that the OpenShift Container Platform installation program generated.
- Your machines have direct internet access or have an HTTP or HTTPS proxy available.

Procedure

1. Monitor the bootstrap process:

```
$ ./openshift-install --dir <installation_directory> wait-for bootstrap-complete \ 1
--log-level=info 2
```

- 1 For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

- 2 To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.

Example output

```
INFO Waiting up to 30m0s for the Kubernetes API at https://api.test.example.com:6443...
INFO API v1.29.4 up
INFO Waiting up to 30m0s for bootstrapping to complete...
INFO It is now safe to remove the bootstrap resources
```

The command succeeds when the Kubernetes API server signals that it has been bootstrapped on the control plane machines.

2. After the bootstrap process is complete, remove the bootstrap machine from the load balancer.



IMPORTANT

You must remove the bootstrap machine from the load balancer at this point. You can also remove or reformat the bootstrap machine itself.

24.1.13. Logging in to the cluster by using the CLI

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

Prerequisites

- You deployed an OpenShift Container Platform cluster.

- You installed the **oc** CLI.

Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

2. Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

Example output

```
system:admin
```

24.1.14. Approving the certificate signing requests for your machines

When you add machines to a cluster, two pending certificate signing requests (CSRs) are generated for each machine that you added. You must confirm that these CSRs are approved or, if necessary, approve them yourself. The client requests must be approved first, followed by the server requests.

Prerequisites

- You added machines to your cluster.

Procedure

1. Confirm that the cluster recognizes the machines:

```
$ oc get nodes
```

Example output

```
NAME      STATUS  ROLES  AGE  VERSION
master-0  Ready   master 63m  v1.29.4
master-1  Ready   master 63m  v1.29.4
master-2  Ready   master 64m  v1.29.4
```

The output lists all of the machines that you created.



NOTE

The preceding output might not include the compute nodes, also known as worker nodes, until some CSRs are approved.

2. Review the pending CSRs and ensure that you see the client requests with the **Pending** or **Approved** status for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

```
NAME      AGE   REQUESTOR                                     CONDITION
csr-8b2br 15m   system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper Pending
csr-8vnps 15m   system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper Pending
...
```

In this example, two machines are joining the cluster. You might see more approved CSRs in the list.

- If the CSRs were not approved, after all of the pending CSRs for the machines you added are in **Pending** status, approve the CSRs for your cluster machines:



NOTE

Because the CSRs rotate automatically, approve your CSRs within an hour of adding the machines to the cluster. If you do not approve them within an hour, the certificates will rotate, and more than two certificates will be present for each node. You must approve all of these certificates. After the client CSR is approved, the Kubelet creates a secondary CSR for the serving certificate, which requires manual approval. Then, subsequent serving certificate renewal requests are automatically approved by the **machine-approver** if the Kubelet requests a new certificate with identical parameters.



NOTE

For clusters running on platforms that are not machine API enabled, such as bare metal and other user-provisioned infrastructure, you must implement a method of automatically approving the kubelet serving certificate requests (CSRs). If a request is not approved, then the **oc exec**, **oc rsh**, and **oc logs** commands cannot succeed, because a serving certificate is required when the API server connects to the kubelet. Any operation that contacts the Kubelet endpoint requires this certificate approval to be in place. The method must watch for new CSRs, confirm that the CSR was submitted by the **node-bootstrapper** service account in the **system:node** or **system:admin** groups, and confirm the identity of the node.

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{\n"}\n{{end}}\n{{end}}' | xargs --no-run-if-empty oc adm certificate approve
```

**NOTE**

Some Operators might not become available until some CSRs are approved.

- Now that your client requests are approved, you must review the server requests for each machine that you added to the cluster:

```
$ oc get csr
```

Example output

```
NAME      AGE    REQUESTOR                                     CONDITION
csr-bfd72 5m26s  system:node:ip-10-0-50-126.us-east-2.compute.internal
Pending
csr-c57lv 5m26s  system:node:ip-10-0-95-157.us-east-2.compute.internal
Pending
...
```

- If the remaining CSRs are not approved, and are in the **Pending** status, approve the CSRs for your cluster machines:

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- 1** **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{"\n"}\n{{end}}{{end}}' | xargs oc adm certificate approve
```

- After all client and server CSRs have been approved, the machines have the **Ready** status. Verify this by running the following command:

```
$ oc get nodes
```

Example output

```
NAME      STATUS  ROLES  AGE  VERSION
master-0  Ready   master 73m  v1.29.4
master-1  Ready   master 73m  v1.29.4
master-2  Ready   master 74m  v1.29.4
worker-0  Ready   worker 11m  v1.29.4
worker-1  Ready   worker 11m  v1.29.4
```

**NOTE**

It can take a few minutes after approval of the server CSRs for the machines to transition to the **Ready** status.

Additional information

- For more information on CSRs, see [Certificate Signing Requests](#).

24.1.15. Initial Operator configuration

After the control plane initializes, you must immediately configure some Operators so that they all become available.

Prerequisites

- Your control plane has initialized.

Procedure

1. Watch the cluster components come online:

```
$ watch -n5 oc get clusteroperators
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.16.0	True	False	False	19m
baremetal	4.16.0	True	False	False	37m
cloud-credential	4.16.0	True	False	False	40m
cluster-autoscaler	4.16.0	True	False	False	37m
config-operator	4.16.0	True	False	False	38m
console	4.16.0	True	False	False	26m
csi-snapshot-controller	4.16.0	True	False	False	37m
dns	4.16.0	True	False	False	37m
etcd	4.16.0	True	False	False	36m
image-registry	4.16.0	True	False	False	31m
ingress	4.16.0	True	False	False	30m
insights	4.16.0	True	False	False	31m
kube-apiserver	4.16.0	True	False	False	26m
kube-controller-manager	4.16.0	True	False	False	36m
kube-scheduler	4.16.0	True	False	False	36m
kube-storage-version-migrator	4.16.0	True	False	False	37m
machine-api	4.16.0	True	False	False	29m
machine-approver	4.16.0	True	False	False	37m
machine-config	4.16.0	True	False	False	36m
marketplace	4.16.0	True	False	False	37m
monitoring	4.16.0	True	False	False	29m
network	4.16.0	True	False	False	38m
node-tuning	4.16.0	True	False	False	37m
openshift-apiserver	4.16.0	True	False	False	32m
openshift-controller-manager	4.16.0	True	False	False	30m
openshift-samples	4.16.0	True	False	False	32m
operator-lifecycle-manager	4.16.0	True	False	False	37m
operator-lifecycle-manager-catalog	4.16.0	True	False	False	37m
operator-lifecycle-manager-packageserver	4.16.0	True	False	False	32m
service-ca	4.16.0	True	False	False	38m
storage	4.16.0	True	False	False	37m

2. Configure the Operators that are not available.

24.1.15.1. Disabling the default OperatorHub catalog sources

Operator catalogs that source content provided by Red Hat and community projects are configured for OperatorHub by default during an OpenShift Container Platform installation. In a restricted network environment, you must disable the default catalogs as a cluster administrator.

Procedure

- Disable the sources for the default catalogs by adding **disableAllDefaultSources: true** to the **OperatorHub** object:

```
$ oc patch OperatorHub cluster --type json \
  -p '[{"op": "add", "path": "/spec/disableAllDefaultSources", "value": true}]'
```

TIP

Alternatively, you can use the web console to manage catalog sources. From the **Administration** → **Cluster Settings** → **Configuration** → **OperatorHub** page, click the **Sources** tab, where you can create, update, delete, disable, and enable individual sources.

24.1.15.2. Image registry removed during installation

On platforms that do not provide shareable object storage, the OpenShift Image Registry Operator bootstraps itself as **Removed**. This allows **openshift-installer** to complete installations on these platform types.

After installation, you must edit the Image Registry Operator configuration to switch the **managementState** from **Removed** to **Managed**. When this has completed, you must configure storage.

24.1.15.3. Image registry storage configuration

The Image Registry Operator is not initially available for platforms that do not provide default storage. After installation, you must configure your registry to use storage so that the Registry Operator is made available.

Instructions are shown for configuring a persistent volume, which is required for production clusters. Where applicable, instructions are shown for configuring an empty directory as the storage location, which is available for only non-production clusters.

Additional instructions are provided for allowing the image registry to use block storage types by using the **Recreate** rollout strategy during upgrades.

24.1.15.3.1. Configuring registry storage for bare metal and other manual installations

As a cluster administrator, following installation you must configure your registry to use storage.

Prerequisites

- You have access to the cluster as a user with the **cluster-admin** role.
- You have a cluster that uses manually-provisioned Red Hat Enterprise Linux CoreOS (RHCOS) nodes, such as bare metal.

- You have provisioned persistent storage for your cluster, such as Red Hat OpenShift Data Foundation.



IMPORTANT

OpenShift Container Platform supports **ReadWriteOnce** access for image registry storage when you have only one replica. **ReadWriteOnce** access also requires that the registry uses the **Recreate** rollout strategy. To deploy an image registry that supports high availability with two or more replicas, **ReadWriteMany** access is required.

- Must have 100Gi capacity.

Procedure

- To configure your registry to use storage, change the **spec.storage.pvc** in the **configs.imageregistry/cluster** resource.



NOTE

When you use shared storage, review your security settings to prevent outside access.

- Verify that you do not have a registry pod:

```
$ oc get pod -n openshift-image-registry -l docker-registry=default
```

Example output

```
No resources found in openshift-image-registry namespace
```



NOTE

If you do have a registry pod in your output, you do not need to continue with this procedure.

- Check the registry configuration:

```
$ oc edit configs.imageregistry.operator.openshift.io
```

Example output

```
storage:
  pvc:
    claim:
```

Leave the **claim** field blank to allow the automatic creation of an **image-registry-storage** PVC.

- Check the **clusteroperator** status:

```
$ oc get clusteroperator image-registry
```


Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
image-registry	4.16	True	False	False	6h50m

5. Ensure that your registry is set to managed to enable building and pushing of images.

- Run:

```
$ oc edit configs.imageregistry/cluster
```

Then, change the line

```
managementState: Removed
```

to

```
managementState: Managed
```

24.1.15.3.2. Configuring storage for the image registry in non-production clusters

You must configure storage for the Image Registry Operator. For non-production clusters, you can set the image registry to an empty directory. If you do so, all images are lost if you restart the registry.

Procedure

- To set the image registry storage to an empty directory:

```
$ oc patch configs.imageregistry.operator.openshift.io cluster --type merge --patch '{"spec": {"storage":{"emptyDir":{}}}'
```



WARNING

Configure this option for only non-production clusters.

If you run this command before the Image Registry Operator initializes its components, the **oc patch** command fails with the following error:

```
Error from server (NotFound): configs.imageregistry.operator.openshift.io "cluster" not found
```

Wait a few minutes and run the command again.

24.1.15.3.3. Configuring block registry storage for bare metal

To allow the image registry to use block storage types during upgrades as a cluster administrator, you can use the **Recreate** rollout strategy.



IMPORTANT

Block storage volumes, or block persistent volumes, are supported but not recommended for use with the image registry on production clusters. An installation where the registry is configured on block storage is not highly available because the registry cannot have more than one replica.

If you choose to use a block storage volume with the image registry, you must use a filesystem persistent volume claim (PVC).

Procedure

1. Enter the following command to set the image registry storage as a block storage type, patch the registry so that it uses the **Recreate** rollout strategy, and runs with only one (**1**) replica:

```
$ oc patch config.imageregistry.operator.openshift.io/cluster --type=merge -p '{"spec": {"rolloutStrategy": "Recreate", "replicas": 1}}'
```

2. Provision the PV for the block storage device, and create a PVC for that volume. The requested block volume uses the ReadWriteOnce (RWO) access mode.
 - a. Create a **pvc.yaml** file with the following contents to define a VMware vSphere **PersistentVolumeClaim** object:

```
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: image-registry-storage 1
  namespace: openshift-image-registry 2
spec:
  accessModes:
  - ReadWriteOnce 3
  resources:
    requests:
      storage: 100Gi 4
```

- 1** A unique name that represents the **PersistentVolumeClaim** object.
- 2** The namespace for the **PersistentVolumeClaim** object, which is **openshift-image-registry**.
- 3** The access mode of the persistent volume claim. With **ReadWriteOnce**, the volume can be mounted with read and write permissions by a single node.
- 4** The size of the persistent volume claim.

- b. Enter the following command to create the **PersistentVolumeClaim** object from the file:

```
$ oc create -f pvc.yaml -n openshift-image-registry
```

3. Enter the following command to edit the registry configuration so that it references the correct PVC:

```
$ oc edit config.imageregistry.operator.openshift.io -o yaml
```

Example output

```
storage:
  pvc:
    claim: ❶
```

- ❶ By creating a custom PVC, you can leave the **claim** field blank for the default automatic creation of an **image-registry-storage** PVC.

24.1.16. Completing installation on user-provisioned infrastructure

After you complete the Operator configuration, you can finish installing the cluster on infrastructure that you provide.

Prerequisites

- Your control plane has initialized.
- You have completed the initial Operator configuration.

Procedure

1. Confirm that all the cluster components are online with the following command:

```
$ watch -n5 oc get clusteroperators
```

Example output

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.16.0	True	False	False	19m
baremetal	4.16.0	True	False	False	37m
cloud-credential	4.16.0	True	False	False	40m
cluster-autoscaler	4.16.0	True	False	False	37m
config-operator	4.16.0	True	False	False	38m
console	4.16.0	True	False	False	26m
csi-snapshot-controller	4.16.0	True	False	False	37m
dns	4.16.0	True	False	False	37m
etcd	4.16.0	True	False	False	36m
image-registry	4.16.0	True	False	False	31m
ingress	4.16.0	True	False	False	30m
insights	4.16.0	True	False	False	31m
kube-apiserver	4.16.0	True	False	False	26m
kube-controller-manager	4.16.0	True	False	False	36m
kube-scheduler	4.16.0	True	False	False	36m
kube-storage-version-migrator	4.16.0	True	False	False	37m
machine-api	4.16.0	True	False	False	29m
machine-approver	4.16.0	True	False	False	37m
machine-config	4.16.0	True	False	False	36m
marketplace	4.16.0	True	False	False	37m
monitoring	4.16.0	True	False	False	29m
network	4.16.0	True	False	False	38m
node-tuning	4.16.0	True	False	False	37m

openshift-apiserver	4.16.0	True	False	False	32m
openshift-controller-manager	4.16.0	True	False	False	30m
openshift-samples	4.16.0	True	False	False	32m
operator-lifecycle-manager	4.16.0	True	False	False	37m
operator-lifecycle-manager-catalog	4.16.0	True	False	False	37m
operator-lifecycle-manager-packageserver	4.16.0	True	False	False	32m
service-ca	4.16.0	True	False	False	38m
storage	4.16.0	True	False	False	37m

Alternatively, the following command notifies you when all of the clusters are available. It also retrieves and displays credentials:

```
$ ./openshift-install --dir <installation_directory> wait-for install-complete 1
```

- 1** For **<installation_directory>**, specify the path to the directory that you stored the installation files in.

Example output

```
INFO Waiting up to 30m0s for the cluster to initialize...
```

The command succeeds when the Cluster Version Operator finishes deploying the OpenShift Container Platform cluster from Kubernetes API server.

IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrapper** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- It is recommended that you use Ignition config files within 12 hours after they are generated because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

2. Confirm that the Kubernetes API server is communicating with the pods.

- a. To view a list of all pods, use the following command:

```
$ oc get pods --all-namespaces
```

Example output

```

NAMESPACE          NAME                                     READY STATUS
RESTARTS  AGE
openshift-apiserver-operator  openshift-apiserver-operator-85cb746d55-zqhs8  1/1
Running      1      9m
```

```

openshift-apiserver      apiserver-67b9g          1/1   Running   0
3m
openshift-apiserver      apiserver-ljcmx          1/1   Running   0
1m
openshift-apiserver      apiserver-z25h4          1/1   Running   0
2m
openshift-authentication-operator authentication-operator-69d5d8bf84-vh2n8 1/1
Running   0      5m
...

```

- b. View the logs for a pod that is listed in the output of the previous command by using the following command:

```
$ oc logs <pod_name> -n <namespace> 1
```

- 1** Specify the pod name and namespace, as shown in the output of the previous command.

If the pod logs display, the Kubernetes API server can communicate with the cluster machines.

3. For an installation with Fibre Channel Protocol (FCP), additional steps are required to enable multipathing. Do not enable multipathing during installation. See "Enabling multipathing with kernel arguments on RHCOS" in the *Postinstallation machine configuration tasks* documentation for more information.

24.1.17. Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.16, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager](#).

After you confirm that your [OpenShift Cluster Manager](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

Additional resources

- See [About remote health monitoring](#) for more information about the Telemetry service

24.1.18. Next steps

- [Customize your cluster](#).
- If necessary, you can [opt out of remote health reporting](#).
- [Set up your registry and configure registry storage](#).

CHAPTER 25. INSTALLATION CONFIGURATION

25.1. CUSTOMIZING NODES

OpenShift Container Platform supports both cluster-wide and per-machine configuration via Ignition, which allows arbitrary partitioning and file content changes to the operating system. In general, if a configuration file is documented in Red Hat Enterprise Linux (RHEL), then modifying it via Ignition is supported.

There are two ways to deploy machine config changes:

- Creating machine configs that are included in manifest files to start up a cluster during **openshift-install**.
- Creating machine configs that are passed to running OpenShift Container Platform nodes via the Machine Config Operator.

Additionally, modifying the reference config, such as the Ignition config that is passed to **coreos-installer** when installing bare-metal nodes allows per-machine configuration. These changes are currently not visible to the Machine Config Operator.

The following sections describe features that you might want to configure on your nodes in this way.

25.1.1. Creating machine configs with Butane

Machine configs are used to configure control plane and worker machines by instructing machines how to create users and file systems, set up the network, install systemd units, and more.

Because modifying machine configs can be difficult, you can use Butane configs to create machine configs for you, thereby making node configuration much easier.

25.1.1.1. About Butane

Butane is a command-line utility that OpenShift Container Platform uses to provide convenient, short-hand syntax for writing machine configs, as well as for performing additional validation of machine configs. The format of the Butane config file that Butane accepts is defined in the [OpenShift Butane config spec](#).

25.1.1.2. Installing Butane

You can install the Butane tool (**butane**) to create OpenShift Container Platform machine configs from a command-line interface. You can install **butane** on Linux, Windows, or macOS by downloading the corresponding binary file.

TIP

Butane releases are backwards-compatible with older releases and with the Fedora CoreOS Config Transpiler (FCCT).

Procedure

1. Navigate to the Butane image download page at <https://mirror.openshift.com/pub/openshift-v4/clients/butane/>.

2. Get the **butane** binary:

- a. For the newest version of Butane, save the latest **butane** image to your current directory:

```
$ curl https://mirror.openshift.com/pub/openshift-v4/clients/butane/latest/butane --output butane
```

- b. Optional: For a specific type of architecture you are installing Butane on, such as aarch64 or ppc64le, indicate the appropriate URL. For example:

```
$ curl https://mirror.openshift.com/pub/openshift-v4/clients/butane/latest/butane-aarch64 --output butane
```

3. Make the downloaded binary file executable:

```
$ chmod +x butane
```

4. Move the **butane** binary file to a directory on your **PATH**.

To check your **PATH**, open a terminal and execute the following command:

```
$ echo $PATH
```

Verification steps

- You can now use the Butane tool by running the **butane** command:

```
$ butane <butane_file>
```

25.1.1.3. Creating a MachineConfig object by using Butane

You can use Butane to produce a **MachineConfig** object so that you can configure worker or control plane nodes at installation time or via the Machine Config Operator.

Prerequisites


- You have installed the **butane** utility.

Procedure

1. Create a Butane config file. The following example creates a file named **99-worker-custom.bu** that configures the system console to show kernel debug messages and specifies custom settings for the chrony time service:

```
variant: openshift
version: 4.16.0
metadata:
  name: 99-worker-custom
  labels:
    machineconfiguration.openshift.io/role: worker
openshift:
  kernel_arguments:
    - loglevel=7
storage:
```

```
files:
- path: /etc/chrony.conf
  mode: 0644
  overwrite: true
  contents:
    inline: |
      pool 0.rhel.pool.ntp.org iburst
      driftfile /var/lib/chrony/drift
      makestep 1.0 3
      rtsync
      logdir /var/log/chrony
```



NOTE

The **99-worker-custom.bu** file is set to create a machine config for worker nodes. To deploy on control plane nodes, change the role from **worker** to **master**. To do both, you could repeat the whole procedure using different file names for the two types of deployments.

2. Create a **MachineConfig** object by giving Butane the file that you created in the previous step:

```
$ butane 99-worker-custom.bu -o ./99-worker-custom.yaml
```

A **MachineConfig** object YAML file is created for you to finish configuring your machines.

3. Save the Butane config in case you need to update the **MachineConfig** object in the future.
4. If the cluster is not running yet, generate manifest files and add the **MachineConfig** object YAML file to the **openshift** directory. If the cluster is already running, apply the file as follows:

```
$ oc create -f 99-worker-custom.yaml
```

Additional resources

- [Adding kernel modules to nodes](#)
- [Encrypting and mirroring disks during installation](#)

25.1.2. Adding day-1 kernel arguments

Although it is often preferable to modify kernel arguments as a day-2 activity, you might want to add kernel arguments to all master or worker nodes during initial cluster installation. Here are some reasons you might want to add kernel arguments during cluster installation so they take effect before the systems first boot up:

- You need to do some low-level network configuration before the systems start.
- You want to disable a feature, such as SELinux, so it has no impact on the systems when they first come up.

**WARNING**

Disabling SELinux on RHCOS in production is not supported. Once SELinux has been disabled on a node, it must be re-provisioned before re-inclusion in a production cluster.

To add kernel arguments to master or worker nodes, you can create a **MachineConfig** object and inject that object into the set of manifest files used by Ignition during cluster setup.

For a listing of arguments you can pass to a RHEL 8 kernel at boot time, see [Kernel.org kernel parameters](#). It is best to only add kernel arguments with this procedure if they are needed to complete the initial OpenShift Container Platform installation.

Procedure

1. Change to the directory that contains the installation program and generate the Kubernetes manifests for the cluster:

```
$ ./openshift-install create manifests --dir <installation_directory>
```

2. Decide if you want to add kernel arguments to worker or control plane nodes.
3. In the **openshift** directory, create a file (for example, **99-openshift-machineconfig-master-kargs.yaml**) to define a **MachineConfig** object to add the kernel settings. This example adds a **loglevel=7** kernel argument to control plane nodes:

```
$ cat << EOF > 99-openshift-machineconfig-master-kargs.yaml
apiVersion: machineconfiguration.openshift.io/v1
kind: MachineConfig
metadata:
  labels:
    machineconfiguration.openshift.io/role: master
  name: 99-openshift-machineconfig-master-kargs
spec:
  kernelArguments:
    - loglevel=7
EOF
```

You can change **master** to **worker** to add kernel arguments to worker nodes instead. Create a separate YAML file to add to both master and worker nodes.

You can now continue on to create the cluster.

25.1.3. Adding kernel modules to nodes

For most common hardware, the Linux kernel includes the device driver modules needed to use that hardware when the computer starts up. For some hardware, however, modules are not available in Linux. Therefore, you must find a way to provide those modules to each host computer. This procedure describes how to do that for nodes in an OpenShift Container Platform cluster.

When a kernel module is first deployed by following these instructions, the module is made available for the current kernel. If a new kernel is installed, the `kmods-via-containers` software will rebuild and deploy the module so a compatible version of that module is available with the new kernel.

The way that this feature is able to keep the module up to date on each node is by:

- Adding a `systemd` service to each node that starts at boot time to detect if a new kernel has been installed and
- If a new kernel is detected, the service rebuilds the module and installs it to the kernel

For information on the software needed for this procedure, see the [kmods-via-containers](#) github site.

A few important issues to keep in mind:

- This procedure is Technology Preview.
- Software tools and examples are not yet available in official RPM form and can only be obtained for now from unofficial **github.com** sites noted in the procedure.
- Third-party kernel modules you might add through these procedures are not supported by Red Hat.
- In this procedure, the software needed to build your kernel modules is deployed in a RHEL 8 container. Keep in mind that modules are rebuilt automatically on each node when that node gets a new kernel. For that reason, each node needs access to a **yum** repository that contains the kernel and related packages needed to rebuild the module. That content is best provided with a valid RHEL subscription.

25.1.3.1. Building and testing the kernel module container

Before deploying kernel modules to your OpenShift Container Platform cluster, you can test the process on a separate RHEL system. Gather the kernel module's source code, the KVC framework, and the `kmod-via-containers` software. Then build and test the module. To do that on a RHEL 8 system, do the following:

Procedure

1. Register a RHEL 8 system:

```
# subscription-manager register
```

2. Attach a subscription to the RHEL 8 system:

```
# subscription-manager attach --auto
```

3. Install software that is required to build the software and container:

```
# yum install podman make git -y
```

4. Clone the **kmod-via-containers** repository:

- a. Create a folder for the repository:

```
$ mkdir kmods; cd kmods
```

- b. Clone the repository:

```
$ git clone https://github.com/kmods-via-containers/kmods-via-containers
```

5. Install a KVC framework instance on your RHEL 8 build host to test the module. This adds a **kmods-via-container** systemd service and loads it:

- a. Change to the **kmod-via-containers** directory:

```
$ cd kmods-via-containers/
```

- b. Install the KVC framework instance:

```
$ sudo make install
```

- c. Reload the systemd manager configuration:

```
$ sudo systemctl daemon-reload
```

6. Get the kernel module source code. The source code might be used to build a third-party module that you do not have control over, but is supplied by others. You will need content similar to the content shown in the **kvc-simple-kmod** example that can be cloned to your system as follows:

```
$ cd .. ; git clone https://github.com/kmods-via-containers/kvc-simple-kmod
```

7. Edit the configuration file, **simple-kmod.conf** file, in this example, and change the name of the Dockerfile to **Dockerfile.rhel**:

- a. Change to the **kvc-simple-kmod** directory:

```
$ cd kvc-simple-kmod
```

- b. Rename the Dockerfile:

```
$ cat simple-kmod.conf
```

Example Dockerfile

```
KMOD_CONTAINER_BUILD_CONTEXT="https://github.com/kmods-via-containers/kvc-
simple-kmod.git"
KMOD_CONTAINER_BUILD_FILE=Dockerfile.rhel
KMOD_SOFTWARE_VERSION=dd1a7d4
KMOD_NAMES="simple-kmod simple-procfs-kmod"
```

8. Create an instance of **kmods-via-containers@.service** for your kernel module, **simple-kmod** in this example:

```
$ sudo make install
```

9. Enable the **kmods-via-containers@.service** instance:

```
$ sudo kmods-via-containers build simple-kmod $(uname -r)
```

10. Enable and start the systemd service:

```
$ sudo systemctl enable kmods-via-containers@simple-kmod.service --now
```

- a. Review the service status:

```
$ sudo systemctl status kmods-via-containers@simple-kmod.service
```

Example output

- `kmods-via-containers@simple-kmod.service - Kmods Via Containers - simple-kmod`
Loaded: loaded (/etc/systemd/system/kmods-via-containers@.service;
enabled; vendor preset: disabled)
Active: active (exited) since Sun 2020-01-12 23:49:49 EST; 5s ago...

11. To confirm that the kernel modules are loaded, use the **lsmod** command to list the modules:

```
$ lsmod | grep simple_
```

Example output

```
simple_procfs_kmod    16384 0
simple_kmod           16384 0
```

12. Optional. Use other methods to check that the **simple-kmod** example is working:

- Look for a "Hello world" message in the kernel ring buffer with **dmesg**:

```
$ dmesg | grep 'Hello world'
```

Example output

```
[ 6420.761332] Hello world from simple_kmod.
```

- Check the value of **simple-procfs-kmod** in **/proc**:

```
$ sudo cat /proc/simple-procfs-kmod
```

Example output

```
simple-procfs-kmod number = 0
```

- Run the **spkut** command to get more information from the module:

```
$ sudo spkut 44
```

Example output

```

KVC: wrapper simple-kmod for 4.18.0-147.3.1.el8_1.x86_64
Running userspace wrapper using the kernel module container...
+ podman run -i --rm --privileged
  simple-kmod-dd1a7d4:4.18.0-147.3.1.el8_1.x86_64 spkut 44
simple-procfs-kmod number = 0
simple-procfs-kmod number = 44

```

Going forward, when the system boots this service will check if a new kernel is running. If there is a new kernel, the service builds a new version of the kernel module and then loads it. If the module is already built, it will just load it.

25.1.3.2. Provisioning a kernel module to OpenShift Container Platform

Depending on whether or not you must have the kernel module in place when OpenShift Container Platform cluster first boots, you can set up the kernel modules to be deployed in one of two ways:

- **Provision kernel modules at cluster install time (day-1)** You can create the content as a **MachineConfig** object and provide it to **openshift-install** by including it with a set of manifest files.
- **Provision kernel modules via Machine Config Operator (day-2)** If you can wait until the cluster is up and running to add your kernel module, you can deploy the kernel module software via the Machine Config Operator (MCO).

In either case, each node needs to be able to get the kernel packages and related software packages at the time that a new kernel is detected. There are a few ways you can set up each node to be able to obtain that content.

- Provide RHEL entitlements to each node.
- Get RHEL entitlements from an existing RHEL host, from the **/etc/pki/entitlement** directory and copy them to the same location as the other files you provide when you build your Ignition config.
- Inside the Dockerfile, add pointers to a **yum** repository containing the kernel and other packages. This must include new kernel packages as they are needed to match newly installed kernels.

25.1.3.2.1. Provision kernel modules via a MachineConfig object

By packaging kernel module software with a **MachineConfig** object, you can deliver that software to worker or control plane nodes at installation time or via the Machine Config Operator.

Procedure

1. Register a RHEL 8 system:

```
# subscription-manager register
```

2. Attach a subscription to the RHEL 8 system:

```
# subscription-manager attach --auto
```

3. Install software needed to build the software:

-

```
# yum install podman make git -y
```

4. Create a directory to host the kernel module and tooling:

```
$ mkdir kmods; cd kmods
```

5. Get the **kmods-via-containers** software:

- a. Clone the **kmods-via-containers** repository:

```
$ git clone https://github.com/kmods-via-containers/kmods-via-containers
```

- b. Clone the **kvc-simple-kmod** repository:

```
$ git clone https://github.com/kmods-via-containers/kvc-simple-kmod
```

6. Get your module software. In this example, **kvc-simple-kmod** is used.

7. Create a fakeroot directory and populate it with files that you want to deliver via Ignition, using the repositories cloned earlier:

- a. Create the directory:

```
$ FAKEROOT=$(mktemp -d)
```

- b. Change to the **kmod-via-containers** directory:

```
$ cd kmods-via-containers
```

- c. Install the KVC framework instance:

```
$ make install DESTDIR=${FAKEROOT}/usr/local CONFDIR=${FAKEROOT}/etc/
```

- d. Change to the **kvc-simple-kmod** directory:

```
$ cd ../kvc-simple-kmod
```

- e. Create the instance:

```
$ make install DESTDIR=${FAKEROOT}/usr/local CONFDIR=${FAKEROOT}/etc/
```

8. Clone the fakeroot directory, replacing any symbolic links with copies of their targets, by running the following command:

```
$ cd .. && rm -rf kmod-tree && cp -Lpr ${FAKEROOT} kmod-tree
```

9. Create a Butane config file, **99-simple-kmod.bu**, that embeds the kernel module tree and enables the systemd service.



NOTE

See "Creating machine configs with Butane" for information about Butane.

```

variant: openshift
version: 4.16.0
metadata:
  name: 99-simple-kmod
  labels:
    machineconfiguration.openshift.io/role: worker 1
storage:
  trees:
    - local: kmod-tree
systemd:
  units:
    - name: kmods-via-containers@simple-kmod.service
      enabled: true

```

- 1** To deploy on control plane nodes, change **worker** to **master**. To deploy on both control plane and worker nodes, perform the remainder of these instructions once for each node type.

10. Use Butane to generate a machine config YAML file, **99-simple-kmod.yaml**, containing the files and configuration to be delivered:

```
$ butane 99-simple-kmod.bu --files-dir . -o 99-simple-kmod.yaml
```

11. If the cluster is not up yet, generate manifest files and add this file to the **openshift** directory. If the cluster is already running, apply the file as follows:

```
$ oc create -f 99-simple-kmod.yaml
```

Your nodes will start the **kmods-via-containers@simple-kmod.service** service and the kernel modules will be loaded.

12. To confirm that the kernel modules are loaded, you can log in to a node (using **oc debug node/<openshift-node>**, then **chroot /host**). To list the modules, use the **lsmod** command:

```
$ lsmod | grep simple_
```

Example output

```

simple_procfs_kmod 16384 0
simple_kmod        16384 0

```

25.1.4. Encrypting and mirroring disks during installation

During an OpenShift Container Platform installation, you can enable boot disk encryption and mirroring on the cluster nodes.

25.1.4.1. About disk encryption

You can enable encryption for the boot disks on the control plane and compute nodes at installation time. OpenShift Container Platform supports the Trusted Platform Module (TPM) v2 and Tang encryption modes.

TPM v2

This is the preferred mode. TPM v2 stores passphrases in a secure cryptoprocessor on the server. You can use this mode to prevent decryption of the boot disk data on a cluster node if the disk is removed from the server.

Tang

Tang and Clevis are server and client components that enable network-bound disk encryption (NBDE). You can bind the boot disk data on your cluster nodes to one or more Tang servers. This prevents decryption of the data unless the nodes are on a secure network where the Tang servers are accessible. Clevis is an automated decryption framework used to implement decryption on the client side.



IMPORTANT

The use of the Tang encryption mode to encrypt your disks is only supported for bare metal and vSphere installations on user-provisioned infrastructure.

In earlier versions of Red Hat Enterprise Linux CoreOS (RHCOS), disk encryption was configured by specifying `/etc/clevis.json` in the Ignition config. That file is not supported in clusters created with OpenShift Container Platform 4.7 or later. Configure disk encryption by using the following procedure.

When the TPM v2 or Tang encryption modes are enabled, the RHCOS boot disks are encrypted using the LUKS2 format.

This feature:

- Is available for installer-provisioned infrastructure, user-provisioned infrastructure, and Assisted Installer deployments
- For Assisted installer deployments:
 - Each cluster can only have a single encryption method, Tang or TPM
 - Encryption can be enabled on some or all nodes
 - There is no Tang threshold; all servers must be valid and operational
 - Encryption applies to the installation disks only, not to the workload disks
- Is supported on Red Hat Enterprise Linux CoreOS (RHCOS) systems only
- Sets up disk encryption during the manifest installation phase, encrypting all data written to disk, from first boot forward
- Requires no user intervention for providing passphrases
- Uses AES-256-XTS encryption, or AES-256-CBC if FIPS mode is enabled

25.1.4.1.1. Configuring an encryption threshold

In OpenShift Container Platform, you can specify a requirement for more than one Tang server. You can also configure the TPM v2 and Tang encryption modes simultaneously. This enables boot disk data decryption only if the TPM secure cryptoprocessor is present and the Tang servers are accessible over a secure network.

You can use the **threshold** attribute in your Butane configuration to define the minimum number of TPM v2 and Tang encryption conditions required for decryption to occur.


The threshold is met when the stated value is reached through any combination of the declared conditions. In the case of offline provisioning, the offline server is accessed using an included advertisement, and only uses that supplied advertisement if the number of online servers do not meet the set threshold.

For example, the **threshold** value of **2** in the following configuration can be reached by accessing two Tang servers, with the offline server available as a backup, or by accessing the TPM secure cryptoprocessor and one of the Tang servers:

Example Butane configuration for disk encryption

```
variant: openshift
version: 4.16.0
metadata:
  name: worker-storage
  labels:
    machineconfiguration.openshift.io/role: worker
boot_device:
  layout: x86_64 1
  luks:
    tpm2: true 2
    tang: 3
      - url: http://tang1.example.com:7500
        thumbprint: jwGN5tRFK-kF6pIX89ssF3khxxX
      - url: http://tang2.example.com:7500
        thumbprint: VCJsvZFjBSIHSlDw78rOrq7h2ZF
      - url: http://tang3.example.com:7500
        thumbprint: PLjNyRdGw03zIRoGjQYMahSZGu9
        advertisement: "{\"payload\": \"...\", \"protected\": \"...\", \"signature\": \"...\"}" 4
    threshold: 2 5
openshift:
  fips: true
```

- 1** Set this field to the instruction set architecture of the cluster nodes. Some examples include, **x86_64**, **aarch64**, or **ppc64le**.
- 2** Include this field if you want to use a Trusted Platform Module (TPM) to encrypt the root file system.
- 3** Include this section if you want to use one or more Tang servers.
- 4** Optional: Include this field for offline provisioning. Ignition will provision the Tang server binding rather than fetching the advertisement from the server at runtime. This lets the server be unavailable at provisioning time.
- 5** Specify the minimum number of TPM v2 and Tang encryption conditions required for decryption to occur.

**IMPORTANT**

The default **threshold** value is **1**. If you include multiple encryption conditions in your configuration but do not specify a threshold, decryption can occur if any of the conditions are met.

**NOTE**

If you require TPM v2 *and* Tang for decryption, the value of the **threshold** attribute must equal the total number of stated Tang servers plus one. If the **threshold** value is lower, it is possible to reach the threshold value by using a single encryption mode. For example, if you set **tpm2** to **true** and specify two Tang servers, a threshold of **2** can be met by accessing the two Tang servers, even if the TPM secure cryptoprocessor is not available.

25.1.4.2. About disk mirroring

During OpenShift Container Platform installation on control plane and worker nodes, you can enable mirroring of the boot and other disks to two or more redundant storage devices. A node continues to function after storage device failure provided one device remains available.

Mirroring does not support replacement of a failed disk. Reprovision the node to restore the mirror to a pristine, non-degraded state.

**NOTE**

For user-provisioned infrastructure deployments, mirroring is available only on RHCOS systems. Support for mirroring is available on **x86_64** nodes booted with BIOS or UEFI and on **ppc64le** nodes.

25.1.4.3. Configuring disk encryption and mirroring

You can enable and configure encryption and mirroring during an OpenShift Container Platform installation.

Prerequisites

- You have downloaded the OpenShift Container Platform installation program on your installation node.
- You installed Butane on your installation node.

**NOTE**

Butane is a command-line utility that OpenShift Container Platform uses to offer convenient, short-hand syntax for writing and validating machine configs. For more information, see "Creating machine configs with Butane".

- You have access to a Red Hat Enterprise Linux (RHEL) 8 machine that can be used to generate a thumbprint of the Tang exchange key.

Procedure

1. If you want to use TPM v2 to encrypt your cluster, check to see if TPM v2 encryption needs to be enabled in the host firmware for each node. This is required on most Dell systems. Check the manual for your specific system.
2. If you want to use Tang to encrypt your cluster, follow these preparatory steps:
 - a. Set up a Tang server or access an existing one. See [Network-bound disk encryption](#) for instructions.
 - b. Install the **clevis** package on a RHEL 8 machine, if it is not already installed:

```
$ sudo yum install clevis
```

- c. On the RHEL 8 machine, run the following command to generate a thumbprint of the exchange key. Replace **http://tang1.example.com:7500** with the URL of your Tang server:

```
$ clevis-encrypt-tang '{"url":"http://tang1.example.com:7500"}' < /dev/null > /dev/null 1
```

- 1** In this example, **tangd.socket** is listening on port **7500** on the Tang server.



NOTE

The **clevis-encrypt-tang** command generates a thumbprint of the exchange key. No data passes to the encryption command during this step; **/dev/null** exists here as an input instead of plain text. The encrypted output is also sent to **/dev/null**, because it is not required for this procedure.

Example output

```
The advertisement contains the following signing keys:
```

```
PLjNyRdGw03zIRoGjQYMahSZGu9 1
```

- 1** The thumbprint of the exchange key.

When the **Do you wish to trust these keys? [ynYN]** prompt displays, type **Y**.

- d. Optional: For offline Tang provisioning:
 - i. Obtain the advertisement from the server using the **curl** command. Replace **http://tang2.example.com:7500** with the URL of your Tang server:

```
$ curl -f http://tang2.example.com:7500/adv > adv.jws && cat adv.jws
```

Expected output

```
{"payload": "eyJrZXlzljogW3siYWxnIjogIktV", "protected":  
"eyJhbGciOiJIUzUxMiIsImN0eSI": "ADLgk7fZdE3Yt4FyYsm0pHiau7Q"}
```

- ii. Provide the advertisement file to Clevis for encryption:

```
$ clevis-encrypt-tang '{"url":"http://tang2.example.com:7500","adv":"adv.jws"}' < /dev/null > /dev/null
```

- e. If the nodes are configured with static IP addressing, run **coreos-installer iso customize --dest-karg-append** or use the **coreos-installer --append-karg** option when installing RHCOS nodes to set the IP address of the installed system. Append the **ip=** and other arguments needed for your network.



IMPORTANT

Some methods for configuring static IPs do not affect the initramfs after the first boot and will not work with Tang encryption. These include the **coreos-installer --copy-network** option, the **coreos-installer iso customize --network-keyfile** option, and the **coreos-installer pxe customize --network-keyfile** option, as well as adding **ip=** arguments to the kernel command line of the live ISO or PXE image during installation. Incorrect static IP configuration causes the second boot of the node to fail.

3. On your installation node, change to the directory that contains the installation program and generate the Kubernetes manifests for the cluster:

```
$. /openshift-install create manifests --dir <installation_directory> 1
```

- 1 Replace **<installation_directory>** with the path to the directory that you want to store the installation files in.

4. Create a Butane config that configures disk encryption, mirroring, or both. For example, to configure storage for compute nodes, create a **\$HOME/clusterconfig/worker-storage.bu** file.

Butane config example for a boot device


```
variant: openshift
version: 4.16.0
metadata:
  name: worker-storage 1
  labels:
    machineconfiguration.openshift.io/role: worker 2
boot_device:
  layout: x86_64 3
  luks: 4
  tpm2: true 5
  tang: 6
    - url: http://tang1.example.com:7500 7
      thumbprint: PLjNyRdGw03zIRoGjQYMahSZGu9 8
    - url: http://tang2.example.com:7500
      thumbprint: VCJsvZFjBSIHSldw78rOrq7h2ZF
      advertisement: '{"payload": "eyJrZXlzljogW3siYWxnIjogIkV", "protected": "eyJhbGciOiJIJFUiLCJ0eXkiOiJ0eXkiLCJ0eXkiOiJ0eXki"}' 9
  threshold: 1 10
  mirror: 11
  devices: 12
```

```

- /dev/sda
- /dev/sdb
openshift:
  fips: true 13

```

- 1** **2** For control plane configurations, replace **worker** with **master** in both of these locations.
- 3** Set this field to the instruction set architecture of the cluster nodes. Some examples include, **x86_64**, **aarch64**, or **ppc64le**.
- 4** Include this section if you want to encrypt the root file system. For more details, see "About disk encryption".
- 5** Include this field if you want to use a Trusted Platform Module (TPM) to encrypt the root file system.
- 6** Include this section if you want to use one or more Tang servers.
- 7** Specify the URL of a Tang server. In this example, **tangd.socket** is listening on port **7500** on the Tang server.
- 8** Specify the exchange key thumbprint, which was generated in a preceding step.
- 9** Optional: Specify the advertisement for your offline Tang server in valid JSON format.
- 10** Specify the minimum number of TPM v2 and Tang encryption conditions that must be met for decryption to occur. The default value is **1**. For more information about this topic, see "Configuring an encryption threshold".
- 11** Include this section if you want to mirror the boot disk. For more details, see "About disk mirroring".
- 12** List all disk devices that should be included in the boot disk mirror, including the disk that RHCOS will be installed onto.
- 13** Include this directive to enable FIPS mode on your cluster.



IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#). If you are configuring nodes to use both disk encryption and mirroring, both features must be configured in the same Butane configuration file. If you are configuring disk encryption on a node with FIPS mode enabled, you must include the **fips** directive in the same Butane configuration file, even if FIPS mode is also enabled in a separate manifest.

5. Create a control plane or compute node manifest from the corresponding Butane configuration file and save it to the **<installation_directory>/openshift** directory. For example, to create a manifest for the compute nodes, run the following command:

```

$ butane $HOME/clusterconfig/worker-storage.bu -o <installation_directory>/openshift/99-
worker-storage.yaml

```

- Repeat this step for each node type that requires disk encryption or mirroring.
- Save the Butane configuration file in case you need to update the manifests in the future.
 - Continue with the remainder of the OpenShift Container Platform installation.

TIP

You can monitor the console log on the RHCOS nodes during installation for error messages relating to disk encryption or mirroring.



IMPORTANT

If you configure additional data partitions, they will not be encrypted unless encryption is explicitly requested.

Verification

After installing OpenShift Container Platform, you can verify if boot disk encryption or mirroring is enabled on the cluster nodes.

- From the installation host, access a cluster node by using a debug pod:
 - Start a debug pod for the node, for example:

```
$ oc debug node/compute-1
```

- Set **/host** as the root directory within the debug shell. The debug pod mounts the root file system of the node in **/host** within the pod. By changing the root directory to **/host**, you can run binaries contained in the executable paths on the node:

```
# chroot /host
```



NOTE

OpenShift Container Platform cluster nodes running Red Hat Enterprise Linux CoreOS (RHCOS) are immutable and rely on Operators to apply cluster changes. Accessing cluster nodes using SSH is not recommended. However, if the OpenShift Container Platform API is not available, or **kubelet** is not properly functioning on the target node, **oc** operations will be impacted. In such situations, it is possible to access nodes using **ssh core@<node>.<cluster_name>.<base_domain>** instead.

- If you configured boot disk encryption, verify if it is enabled:
 - From the debug shell, review the status of the root mapping on the node:

```
# cryptsetup status root
```

Example output

```
/dev/mapper/root is active and is in use.  
type: LUKS2 1
```

```

cipher: aes-xts-plain64 2
keysize: 512 bits
key location: keyring
device: /dev/sda4 3
sector size: 512
offset: 32768 sectors
size: 15683456 sectors
mode: read/write

```

- 1 The encryption format. When the TPM v2 or Tang encryption modes are enabled, the RHCOS boot disks are encrypted using the LUKS2 format.
- 2 The encryption algorithm used to encrypt the LUKS2 volume. The **aes-cbc-essiv:sha256** cipher is used if FIPS mode is enabled.
- 3 The device that contains the encrypted LUKS2 volume. If mirroring is enabled, the value will represent a software mirror device, for example **/dev/md126**.

b. List the Clevis plugins that are bound to the encrypted device:

```
# clevis luks list -d /dev/sda4 1
```

- 1 Specify the device that is listed in the **device** field in the output of the preceding step.

Example output

```
1: sss '{"t":1,"pins":{"tang":{"url":"http://tang.example.com:7500"}}}' 1
```

- 1 In the example output, the Tang plugin is used by the Shamir's Secret Sharing (SSS) Clevis plugin for the **/dev/sda4** device.

3. If you configured mirroring, verify if it is enabled:

a. From the debug shell, list the software RAID devices on the node:

```
# cat /proc/mdstat
```

Example output

```

Personalities : [raid1]
md126 : active raid1 sdb3[1] sda3[0] 1
      393152 blocks super 1.0 [2/2] [UU]

md127 : active raid1 sda4[0] sdb4[1] 2
      51869632 blocks super 1.2 [2/2] [UU]

unused devices: <none>

```

- 1 The **/dev/md126** software RAID mirror device uses the **/dev/sda3** and **/dev/sdb3** disk devices on the cluster node.

- 2 The `/dev/md127` software RAID mirror device uses the `/dev/sda4` and `/dev/sdb4` disk devices on the cluster node.

- b. Review the details of each of the software RAID devices listed in the output of the preceding command. The following example lists the details of the `/dev/md126` device:

```
# mdadm --detail /dev/md126
```

Example output

```
/dev/md126:
  Version : 1.0
  Creation Time : Wed Jul 7 11:07:36 2021
  Raid Level : raid1 1
  Array Size : 393152 (383.94 MiB 402.59 MB)
  Used Dev Size : 393152 (383.94 MiB 402.59 MB)
  Raid Devices : 2
  Total Devices : 2
  Persistence : Superblock is persistent

  Update Time : Wed Jul 7 11:18:24 2021
  State : clean 2
  Active Devices : 2 3
  Working Devices : 2 4
  Failed Devices : 0 5
  Spare Devices : 0

Consistency Policy : resync

  Name : any:md-boot 6
  UUID : cdfa3801:c520e0b5:2bee2755:69043055
  Events : 19

  Number Major Minor RaidDevice State
    0   252    3    0   active sync  /dev/sda3 7
    1   252   19    1   active sync  /dev/sdb3 8
```

- 1 Specifies the RAID level of the device. **raid1** indicates RAID 1 disk mirroring.
- 2 Specifies the state of the RAID device.
- 3 4 States the number of underlying disk devices that are active and working.
- 5 States the number of underlying disk devices that are in a failed state.
- 6 The name of the software RAID device.
- 7 8 Provides information about the underlying disk devices used by the software RAID device.

- c. List the file systems mounted on the software RAID devices:


```
# mount | grep /dev/md
```

Example output

```
/dev/md127 on / type xfs
(rw,relatime,seclabel,attr2,inode64,logbufs=8,logbsize=32k,prjquota)
/dev/md127 on /etc type xfs
(rw,relatime,seclabel,attr2,inode64,logbufs=8,logbsize=32k,prjquota)
/dev/md127 on /usr type xfs
(ro,relatime,seclabel,attr2,inode64,logbufs=8,logbsize=32k,prjquota)
/dev/md127 on /sysroot type xfs
(ro,relatime,seclabel,attr2,inode64,logbufs=8,logbsize=32k,prjquota)
/dev/md127 on /var type xfs
(rw,relatime,seclabel,attr2,inode64,logbufs=8,logbsize=32k,prjquota)
/dev/md127 on /var/lib/containers/storage/overlay type xfs
(rw,relatime,seclabel,attr2,inode64,logbufs=8,logbsize=32k,prjquota)
/dev/md127 on /var/lib/kubelet/pods/e5054ed5-f882-4d14-b599-99c050d4e0c0/volume-
subpaths/etc/tuned/1 type xfs
(rw,relatime,seclabel,attr2,inode64,logbufs=8,logbsize=32k,prjquota)
/dev/md127 on /var/lib/kubelet/pods/e5054ed5-f882-4d14-b599-99c050d4e0c0/volume-
subpaths/etc/tuned/2 type xfs
(rw,relatime,seclabel,attr2,inode64,logbufs=8,logbsize=32k,prjquota)
/dev/md127 on /var/lib/kubelet/pods/e5054ed5-f882-4d14-b599-99c050d4e0c0/volume-
subpaths/etc/tuned/3 type xfs
(rw,relatime,seclabel,attr2,inode64,logbufs=8,logbsize=32k,prjquota)
/dev/md127 on /var/lib/kubelet/pods/e5054ed5-f882-4d14-b599-99c050d4e0c0/volume-
subpaths/etc/tuned/4 type xfs
(rw,relatime,seclabel,attr2,inode64,logbufs=8,logbsize=32k,prjquota)
/dev/md127 on /var/lib/kubelet/pods/e5054ed5-f882-4d14-b599-99c050d4e0c0/volume-
subpaths/etc/tuned/5 type xfs
(rw,relatime,seclabel,attr2,inode64,logbufs=8,logbsize=32k,prjquota)
/dev/md126 on /boot type ext4 (rw,relatime,seclabel)
```

In the example output, the **/boot** file system is mounted on the **/dev/md126** software RAID device and the root file system is mounted on **/dev/md127**.

4. Repeat the verification steps for each OpenShift Container Platform node type.

Additional resources

- For more information about the TPM v2 and Tang encryption modes, see [Configuring automated unlocking of encrypted volumes using policy-based decryption](#).

25.1.4.4. Configuring a RAID-enabled data volume

You can enable software RAID partitioning to provide an external data volume. OpenShift Container Platform supports RAID 0, RAID 1, RAID 4, RAID 5, RAID 6, and RAID 10 for data protection and fault tolerance. See "About disk mirroring" for more details.

Prerequisites

- You have downloaded the OpenShift Container Platform installation program on your installation node.

- You have installed Butane on your installation node.



NOTE

Butane is a command-line utility that OpenShift Container Platform uses to provide convenient, short-hand syntax for writing machine configs, as well as for performing additional validation of machine configs. For more information, see the *Creating machine configs with Butane* section.

Procedure

- Create a Butane config that configures a data volume by using software RAID.
 - To configure a data volume with RAID 1 on the same disks that are used for a mirrored boot disk, create a `$HOME/clusterconfig/raid1-storage.bu` file, for example:

RAID 1 on mirrored boot disk

```
variant: openshift
version: 4.16.0
metadata:
  name: raid1-storage
  labels:
    machineconfiguration.openshift.io/role: worker
boot_device:
  mirror:
    devices:
      - /dev/disk/by-id/scsi-3600508b400105e210000900000490000
      - /dev/disk/by-id/scsi-SSEAGATE_ST373453LW_3HW1RHM6
storage:
  disks:
    - device: /dev/disk/by-id/scsi-3600508b400105e210000900000490000
      partitions:
        - label: root-1
          size_mib: 25000 1
        - label: var-1
    - device: /dev/disk/by-id/scsi-SSEAGATE_ST373453LW_3HW1RHM6
      partitions:
        - label: root-2
          size_mib: 25000 2
        - label: var-2
  raid:
    - name: md-var
      level: raid1
      devices:
        - /dev/disk/by-partlabel/var-1
        - /dev/disk/by-partlabel/var-2
  filesystems:
    - device: /dev/md/md-var
      path: /var
      format: xfs
      wipe_filesystem: true
      with_mount_unit: true
```

1 2 When adding a data partition to the boot disk, a minimum value of 25000 mebibytes is recommended. If no value is specified, or if the specified value is smaller than the

- To configure a data volume with RAID 1 on secondary disks, create a **\$HOME/clusterconfig/raid1-alt-storage.bu** file, for example:

RAID 1 on secondary disks

```
variant: openshift
version: 4.16.0
metadata:
  name: raid1-alt-storage
  labels:
    machineconfiguration.openshift.io/role: worker
storage:
  disks:
    - device: /dev/sdc
      wipe_table: true
      partitions:
        - label: data-1
    - device: /dev/sdd
      wipe_table: true
      partitions:
        - label: data-2
  raid:
    - name: md-var-lib-containers
      level: raid1
      devices:
        - /dev/disk/by-partlabel/data-1
        - /dev/disk/by-partlabel/data-2
  filesystems:
    - device: /dev/md/md-var-lib-containers
      path: /var/lib/containers
      format: xfs
      wipe_filesystem: true
      with_mount_unit: true
```

2. Create a RAID manifest from the Butane config you created in the previous step and save it to the **<installation_directory>/openshift** directory. For example, to create a manifest for the compute nodes, run the following command:

```
$ butane $HOME/clusterconfig/<butane_config>.bu -o
<installation_directory>/openshift/<manifest_name>.yaml 1
```

- 1** Replace **<butane_config>** and **<manifest_name>** with the file names from the previous step. For example, **raid1-alt-storage.bu** and **raid1-alt-storage.yaml** for secondary disks.

3. Save the Butane config in case you need to update the manifest in the future.
4. Continue with the remainder of the OpenShift Container Platform installation.

25.1.4.5. Configuring an Intel® Virtual RAID on CPU (VROC) data volume

Intel® VROC is a type of hybrid RAID, where some of the maintenance is offloaded to the hardware, but appears as software RAID to the operating system.

The following procedure configures an Intel® VROC-enabled RAID1.

Prerequisites

- You have a system with Intel® Volume Management Device (VMD) enabled.

Procedure

1. Create the Intel® Matrix Storage Manager (IMSM) RAID container by running the following command:

```
$ mdadm -CR /dev/md/ism0 -e \  
ism -n2 /dev/nvme0n1 /dev/nvme1n1 1
```

- 1** The RAID device names. In this example, there are two devices listed. If you provide more than two device names, you must adjust the **-n** flag. For example, listing three devices would use the flag **-n3**.

2. Create the RAID1 storage inside the container:

- a. Create a dummy RAID0 volume in front of the real RAID1 volume by running the following command:

```
$ mdadm -CR /dev/md/dummy -l0 -n2 /dev/ism0 -z10m --assume-clean
```

- b. Create the real RAID1 array by running the following command:

```
$ mdadm -CR /dev/md/coreos -l1 -n2 /dev/ism0
```

- c. Stop both RAID0 and RAID1 member arrays and delete the dummy RAID0 array with the following commands:

```
$ mdadm -S /dev/md/dummy \  
mdadm -S /dev/md/coreos \  
mdadm --kill-subarray=0 /dev/md/ism0
```

- d. Restart the RAID1 arrays by running the following command:

```
$ mdadm -A /dev/md/coreos /dev/md/ism0
```

3. Install RHCOS on the RAID1 device:

- a. Get the UUID of the IMSM container by running the following command:

```
$ mdadm --details --export /dev/md/ism0
```

- b. Install RHCOS and include the **rd.md.uuid** kernel argument by running the following command:

```
$ coreos-installer install /dev/md/coreos \
  --append-karg rd.md.uuid=<md_UUID> ❶
...
```

- ❶ The UUID of the IMSM container.

Include any additional **coreos-installer** arguments you need to install RHCOS.

25.1.5. Configuring chrony time service

You can set the time server and related settings used by the chrony time service (**chronyd**) by modifying the contents of the **chrony.conf** file and passing those contents to your nodes as a machine config.

Procedure

1. Create a Butane config including the contents of the **chrony.conf** file. For example, to configure chrony on worker nodes, create a **99-worker-chrony.bu** file.



NOTE

See "Creating machine configs with Butane" for information about Butane.

```
variant: openshift
version: 4.16.0
metadata:
  name: 99-worker-chrony ❶
labels:
  machineconfiguration.openshift.io/role: worker ❷
storage:
  files:
  - path: /etc/chrony.conf
    mode: 0644 ❸
    overwrite: true
    contents:
      inline: |
        pool 0.rhel.pool.ntp.org iburst ❹
        driftfile /var/lib/chrony/drift
        makestep 1.0 3
        rtsync
        logdir /var/log/chrony
```

- ❶ ❷ On control plane nodes, substitute **master** for **worker** in both of these locations.
- ❸ Specify an octal value mode for the **mode** field in the machine config file. After creating the file and applying the changes, the **mode** is converted to a decimal value. You can check the YAML file with the command **oc get mc <mc-name> -o yaml**.
- ❹ Specify any valid, reachable time source, such as the one provided by your DHCP server. Alternately, you can specify any of the following NTP servers: **1.rhel.pool.ntp.org**, **2.rhel.pool.ntp.org**, or **3.rhel.pool.ntp.org**.

- Use Butane to generate a **MachineConfig** object file, **99-worker-chrony.yaml**, containing the configuration to be delivered to the nodes:

```
$ butane 99-worker-chrony.bu -o 99-worker-chrony.yaml
```

- Apply the configurations in one of two ways:

- If the cluster is not running yet, after you generate manifest files, add the **MachineConfig** object file to the **<installation_directory>/openshift** directory, and then continue to create the cluster.
- If the cluster is already running, apply the file:

```
$ oc apply -f ./99-worker-chrony.yaml
```

25.1.6. Additional resources

- For information on Butane, see [Creating machine configs with Butane](#).
- For information on FIPS support, see [Support for FIPS cryptography](#).


25.2. CONFIGURING YOUR FIREWALL

If you use a firewall, you must configure it so that OpenShift Container Platform can access the sites that it requires to function. You must always grant access to some sites, and you grant access to more if you use Red Hat Insights, the Telemetry service, a cloud to host your cluster, and certain build strategies.

25.2.1. Configuring your firewall for OpenShift Container Platform

Before you install OpenShift Container Platform, you must configure your firewall to grant access to the sites that OpenShift Container Platform requires. When using a firewall, make additional configurations to the firewall so that OpenShift Container Platform can access the sites that it requires to function.

There are no special configuration considerations for services running on only controller nodes compared to worker nodes.



NOTE

If your environment has a dedicated load balancer in front of your OpenShift Container Platform cluster, review the allowlists between your firewall and load balancer to prevent unwanted network restrictions to your cluster.

Procedure

- Set the following registry URLs for your firewall's allowlist:

URL	Port	Function
registry.redhat.io	443	Provides core container images

URL	Port	Function
access.redhat.com ^[1]	443	Hosts all the container images that are stored on the Red Hat Ecosystem Catalog, including core container images.
quay.io	443	Provides core container images
cdn.quay.io	443	Provides core container images
cdn01.quay.io	443	Provides core container images
cdn02.quay.io	443	Provides core container images
cdn03.quay.io	443	Provides core container images
cdn04.quay.io	443	Provides core container images
cdn05.quay.io	443	Provides core container images
cdn06.quay.io	443	Provides core container images
sso.redhat.com	443	The https://console.redhat.com site uses authentication from sso.redhat.com

1. In a firewall environment, ensure that the **access.redhat.com** resource is on the allowlist. This resource hosts a signature store that a container client requires for verifying images when pulling them from **registry.access.redhat.com**.

You can use the wildcards ***.quay.io** and ***.openshiftapps.com** instead of **cdn.quay.io** and **cdn0[1-3].quay.io** in your allowlist. When you add a site, such as **quay.io**, to your allowlist, do not add a wildcard entry, such as ***.quay.io**, to your denylist. In most cases, image registries use a content delivery network (CDN) to serve images. If a firewall blocks access, image downloads are denied when the initial download request redirects to a hostname such as **cdn01.quay.io**.

2. Set your firewall's allowlist to include any site that provides resources for a language or framework that your builds require.
3. If you do not disable Telemetry, you must grant access to the following URLs to access Red Hat Insights:

URL	Port	Function
cert-api.access.redhat.com	443	Required for Telemetry
api.access.redhat.com	443	Required for Telemetry
infogw.api.openshift.com	443	Required for Telemetry

URL	Port	Function
console.redhat.com	443	Required for Telemetry and for insights-operator

4. If you use Alibaba Cloud, Amazon Web Services (AWS), Microsoft Azure, or Google Cloud Platform (GCP) to host your cluster, you must grant access to the URLs that offer the cloud provider API and DNS for that cloud:

Cloud	URL	Port	Function
Alibaba	*.aliyuncs.com	443	Required to access Alibaba Cloud services and resources. Review the Alibaba endpoints_config.go file to find the exact endpoints to allow for the regions that you use.
AWS	aws.amazon.com	443	Used to install and manage clusters in an AWS environment.
	*.amazonaws.com Alternatively, if you choose to not use a wildcard for AWS APIs, you must include the following URLs in your allowlist:	443	Required to access AWS services and resources. Review the AWS Service Endpoints in the AWS documentation to find the exact endpoints to allow for the regions that you use.
	ec2.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	events.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	iam.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	route53.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	*.s3.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	*.s3.<aws_region>.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	*.s3.dualstack.<aws_region>.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
sts.amazonaws.com	443	Used to install and manage clusters in an AWS environment.	

Cloud	URL	Port	Function
	sts. <aws_region>.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	tagging.us-east-1. amazonaws.com	443	Used to install and manage clusters in an AWS environment. This endpoint is always us-east-1 , regardless of the region the cluster is deployed in.
	ec2. <aws_region>.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	elasticloadbalancing. <aws_region>.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	servicequotas. <aws_region>.amazonaws.com	443	Required. Used to confirm quotas for deploying the service.
	tagging. <aws_region>.amazonaws.com	443	Allows the assignment of metadata about AWS resources in the form of tags.
	*.cloudfront.net	443	Used to provide access to CloudFront. If you use the AWS Security Token Service (STS) and the private S3 bucket, you must provide access to CloudFront.
GCP	*.googleapis.com	443	Required to access GCP services and resources. Review Cloud Endpoints in the GCP documentation to find the endpoints to allow for your APIs.
	accounts.google.com	443	Required to access your GCP account.
Microsoft Azure	management.azure.com	443	Required to access Microsoft Azure services and resources. Review the Microsoft Azure REST API reference in the Microsoft Azure documentation to find the endpoints to allow for your APIs.
	*.blob.core.windows.net	443	Required to download Ignition files.

Cloud	URL	Port	Function
	login.microsoftonline.com	443	Required to access Microsoft Azure services and resources. Review the Azure REST API reference in the Microsoft Azure documentation to find the endpoints to allow for your APIs.

5. Allowlist the following URLs:

URL	Port	Function
mirror.openshift.com	443	Required to access mirrored installation content and images. This site is also a source of release image signatures, although the Cluster Version Operator needs only a single functioning source.
storage.googleapis.com/openshift-release	443	A source of release image signatures, although the Cluster Version Operator needs only a single functioning source.
*.apps.<cluster_name>.<base_domain>	443	Required to access the default cluster routes unless you set an ingress wildcard during installation.
quayio-production-s3.s3.amazonaws.com	443	Required to access Quay image content in AWS.
api.openshift.com	443	Required both for your cluster token and to check if updates are available for the cluster.
rhcos.mirror.openshift.com	443	Required to download Red Hat Enterprise Linux CoreOS (RHCOS) images.
console.redhat.com	443	Required for your cluster token.
sso.redhat.com	443	The https://console.redhat.com site uses authentication from sso.redhat.com

Operators require route access to perform health checks. Specifically, the authentication and web console Operators connect to two routes to verify that the routes work. If you are the cluster administrator and do not want to allow ***.apps.<cluster_name>.<base_domain>**, then allow these routes:

- **oauth-openshift.apps.<cluster_name>.<base_domain>**

- **console-openshift-console.apps.<cluster_name>.<base_domain>**, or the hostname that is specified in the **spec.route.hostname** field of the **consoles.operator/cluster** object if the field is not empty.
6. Allowlist the following URLs for optional third-party content:

URL	Port	Function
registry.connect.redhat.com	443	Required for all third-party images and certified operators.
rhc4tp-prod-z8cxf-image-registry-us-east-1-evenkyleffocxqvofrk.s3.dualstack.us-east-1.amazonaws.com	443	Provides access to container images hosted on registry.connect.redhat.com
oso-rhc4tp-docker-registry.s3-us-west-2.amazonaws.com	443	Required for Sonatype Nexus, F5 Big IP operators.

7. If you use a default Red Hat Network Time Protocol (NTP) server allow the following URLs:

- **1.rhel.pool.ntp.org**
- **2.rhel.pool.ntp.org**
- **3.rhel.pool.ntp.org**



NOTE

If you do not use a default Red Hat NTP server, verify the NTP server for your platform and allow it in your firewall.

Additional resources

- [OpenID Connect requirements for AWS STS](#)

25.2.2. OpenShift Container Platform network flow matrix

The network flow matrix describes the ingress flows to OpenShift Container Platform services. The network information in the matrix is accurate for both bare-metal and cloud environments. Use the information in the network flow matrix to help you manage ingress traffic. You can restrict ingress traffic to essential flows to improve network security.

To view or download the raw CSV content, see [this resource](#).

Additionally, consider the following dynamic port ranges when managing ingress traffic:

- **9000-9999**: Host level services
- **3000-32767**: Kubernetes node ports
- **49152-65535**: Dynamic or private ports

**NOTE**

The network flow matrix describes ingress traffic flows for a base OpenShift Container Platform installation. It does not describe network flows for additional components, such as optional Operators available from the Red Hat Marketplace. The matrix does not apply for Hosted-Control-Plane, MicroShift, or standalone clusters.

Table 25.1. Network flow matrix

Direction	Protocol	Port	Namespace	Service	Pod	Container	Node Role	Optional
Ingress	TCP	22	Host system service	sshd			master	TRUE
Ingress	TCP	53	openshift-dns	dns-default	dnf-default	dns	master	FALSE
Ingress	TCP	111	Host system service	rpcbind			master	TRUE
Ingress	TCP	2379	openshift-etcd	etcd	etcd	etcdctl	master	FALSE
Ingress	TCP	2380	openshift-etcd	healthz	etcd	etcd	master	FALSE
Ingress	TCP	5050	openshift-machine-api		ironic-proxy	ironic-proxy	master	FALSE
Ingress	TCP	6080	openshift-kube-apiserver		kube-apiserver	kube-apiserver-insecure-readyz	master	FALSE
Ingress	TCP	6385	openshift-machine-api		ironic-proxy	ironic-proxy	master	FALSE
Ingress	TCP	6443	openshift-kube-apiserver	apiserver	kube-apiserver	kube-apiserver	master	FALSE

Direction	Protocol	Port	Namespace	Service	Pod	Container	Node Role	Optional
Ingress	TCP	8080	openshift-network-operator		network-operator	network-operator	master	FALSE
Ingress	TCP	8798	openshift-machine-config-operator	machine-config-daemon	machine-config-daemon	machine-config-daemon	master	FALSE
Ingress	TCP	9001	openshift-machine-config-operator	machine-config-daemon	machine-config-daemon	kube-rbac-proxy	master	FALSE
Ingress	TCP	9099	openshift-cluster-version	cluster-version-operator	cluster-version-operator	cluster-version-operator	master	FALSE
Ingress	TCP	9100	openshift-monitoring	node-exporter	node-exporter	kube-rbac-proxy	master	FALSE
Ingress	TCP	9103	openshift-ovn-kubernetes	ovn-kubernetes-node	ovnkube-node	kube-rbac-proxy-node	master	FALSE
Ingress	TCP	9104	openshift-network-operator	metrics	network-operator	network-operator	master	FALSE
Ingress	TCP	9105	openshift-ovn-kubernetes	ovn-kubernetes-node	ovnkube-node	kube-rbac-proxy-ovn-metrics	master	FALSE

Direction	Protocol	Port	Namespace	Service	Pod	Container	Node Role	Optional
Ingress	TCP	9107	openshift-ovn-kubernetes	egressip-node-healthcheck	ovnkube-node	ovnkube-controller	master	FALSE
Ingress	TCP	9108	openshift-ovn-kubernetes	ovn-kubernetes-control-plane	ovnkube-control-plane	kube-rbac-proxy	master	FALSE
Ingress	TCP	9192	openshift-cluster-machine-approver	machine-approver	machine-approver	kube-rbac-proxy	master	FALSE
Ingress	TCP	9258	openshift-cloud-controller-manager-operator	machine-approver	cluster-cloud-controller-manager	cluster-cloud-controller-manager	master	FALSE
Ingress	TCP	9444	openshift-kni-infra		haproxy	haproxy	master	FALSE
Ingress	TCP	9445	openshift-kni-infra		haproxy	haproxy	master	FALSE
Ingress	TCP	9447	openshift-machine-api		metal3-baremetal-operator		master	FALSE
Ingress	TCP	9537	Host system service	crio-metrics			master	FALSE
Ingress	TCP	9637	openshift-machine-config-operator	kube-rbac-proxy-crio	kube-rbac-proxy-crio	kube-rbac-proxy-crio	master	FALSE

Direction	Protocol	Port	Namespace	Service	Pod	Container	Node Role	Optional
Ingress	TCP	9978	openshift-etcd	etcd	etcd	etcd-metrics	master	FALSE
Ingress	TCP	9979	openshift-etcd	etcd	etcd	etcd-metrics	master	FALSE
Ingress	TCP	9980	openshift-etcd	etcd	etcd	etcd	master	FALSE
Ingress	TCP	10250	Host system service	kubelet			master	FALSE
Ingress	TCP	10256	openshift-ovn-kubernetes	ovnkube	ovnkube	ovnkube-controller	master	FALSE
Ingress	TCP	10257	openshift-kube-controller-manager	kube-controller-manager	kube-controller-manager	kube-controller-manager	master	FALSE
Ingress	TCP	10258	openshift-cloud-controller-manager-operator	cloud-controller	cloud-controller-manager	cloud-controller-manager	master	FALSE
Ingress	TCP	10259	openshift-kube-scheduler	scheduler	openshift-kube-scheduler	kube-scheduler	master	FALSE
Ingress	TCP	10260	openshift-cloud-controller-manager-operator	cloud-controller	cloud-controller-manager	cloud-controller-manager	master	FALSE

Direction	Protocol	Port	Namespace	Service	Pod	Container	Node Role	Optional
Ingress	TCP	10300	openshift-cluster-csi-drivers	csi-liveness-probe	csi-driver-node	csi-driver	master	FALSE
Ingress	TCP	10309	openshift-cluster-csi-drivers	csi-node-driver	csi-driver-node	csi-node-driver-registry	master	FALSE
Ingress	TCP	10357	openshift-kube-apiserver	openshift-kube-apiserver-healthz	kube-apiserver	kube-apiserver-check-endpoints	master	FALSE
Ingress	TCP	17697	openshift-kube-apiserver	openshift-kube-apiserver-healthz	kube-apiserver	kube-apiserver-check-endpoints	master	FALSE
Ingress	TCP	18080	openshift-kni-infra		coredns	coredns	master	FALSE
Ingress	TCP	22623	openshift-machine-config-operator	machine-config-server	machine-config-server	machine-config-server	master	FALSE
Ingress	TCP	22624	openshift-machine-config-operator	machine-config-server	machine-config-server	machine-config-server	master	FALSE
Ingress	UDP	53	openshift-dns	dns-default	dnf-default	dns	master	FALSE
Ingress	UDP	111	Host system service	rpcbind			master	TRUE


Direction	Protocol	Port	Namespace	Service	Pod	Container	Node Role	Optional
Ingress	UDP	6081	openshift-ovn-kubernetes	ovn-kubernetes-geneve			master	FALSE
Ingress	TCP	22	Host system service	sshd			worker	TRUE
Ingress	TCP	53	openshift-dns	dns-default	dnf-default	dns	worker	FALSE
Ingress	TCP	80	openshift-ingress	router-default	router-default	router	worker	FALSE
Ingress	TCP	111	Host system service	rpcbind			worker	TRUE
Ingress	TCP	443	openshift-ingress	router-default	router-default	router	worker	FALSE
Ingress	TCP	8798	openshift-machine-config-operator	machine-config-daemon	machine-config-daemon	machine-config-daemon	worker	FALSE
Ingress	TCP	9001	openshift-machine-config-operator	machine-config-daemon	machine-config-daemon	kube-rbac-proxy	worker	FALSE
Ingress	TCP	9100	openshift-monitoring	node-exporter	node-exporter	kube-rbac-proxy	worker	FALSE
Ingress	TCP	9103	openshift-ovn-kubernetes	ovn-kubernetes-node	ovnkube-node	kube-rbac-proxy-node	worker	FALSE

Direction	Protocol	Port	Namespace	Service	Pod	Container	Node Role	Optional
Ingress	TCP	9105	openshift-ovn-kubernetes	ovn-kubernetes-node	ovnkube-node	kube-rbac-proxy-ovn-metrics	worker	FALSE
Ingress	TCP	9107	openshift-ovn-kubernetes	egressip-node-healthcheck	ovnkube-node	ovnkube-controller	worker	FALSE
Ingress	TCP	9537	Host system service	crio-metrics			worker	FALSE
Ingress	TCP	9637	openshift-machine-config-operator	kube-rbac-proxy-crio	kube-rbac-proxy-crio	kube-rbac-proxy-crio	worker	FALSE
Ingress	TCP	10250	Host system service	kubelet			worker	FALSE
Ingress	TCP	10256	openshift-ovn-kubernetes	ovnkube	ovnkube	ovnkube-controller	worker	TRUE
Ingress	TCP	10300	openshift-cluster-csi-drivers	csi-liveness-probe	csi-driver-node	csi-driver	worker	FALSE
Ingress	TCP	10309	openshift-cluster-csi-drivers	csi-node-driver-registrar	csi-driver-node	csi-node-driver-registrar	worker	FALSE
Ingress	TCP	18080	openshift-kni-infra		coredns	coredns	worker	FALSE

Direction	Protocol	Port	Namespace	Service	Pod	Container	Node Role	Optional
Ingress	UDP	53	openshift-dns	dns-default	dnf-default	dns	worker	FALSE
Ingress	UDP	111	Host system service	rpcbind			worker	TRUE
Ingress	UDP	6081	openshift-ovn-kubernetes	ovn-kubernetes-geneve			worker	FALSE

25.3. ENABLING LINUX CONTROL GROUP VERSION 1 (CGROUP V1)

As of OpenShift Container Platform 4.14, OpenShift Container Platform uses [Linux control group version 2](#) (cgroup v2) in your cluster. If you are using cgroup v1 on OpenShift Container Platform 4.13 or earlier, migrating to OpenShift Container Platform 4.16 will not automatically update your cgroup configuration to version 2. A fresh installation of OpenShift Container Platform 4.14 or later will use cgroup v2 by default. However, you can enable [Linux control group version 1](#) (cgroup v1) upon installation. Enabling cgroup v1 in OpenShift Container Platform disables all cgroup v2 controllers and hierarchies in your cluster.



IMPORTANT

cgroup v1 is a deprecated feature. Deprecated functionality is still included in OpenShift Container Platform and continues to be supported; however, it will be removed in a future release of this product and is not recommended for new deployments.

For the most recent list of major functionality that has been deprecated or removed within OpenShift Container Platform, refer to the *Deprecated and removed features* section of the OpenShift Container Platform release notes.

cgroup v2 is the current version of the Linux cgroup API. cgroup v2 offers several improvements over cgroup v1, including a unified hierarchy, safer sub-tree delegation, new features such as [Pressure Stall Information](#), and enhanced resource management and isolation. However, cgroup v2 has different CPU, memory, and I/O management characteristics than cgroup v1. Therefore, some workloads might experience slight differences in memory or CPU usage on clusters that run cgroup v2.

You can switch between cgroup v1 and cgroup v2, as needed, by editing the **node.config** object. For more information, see "Configuring the Linux cgroup on your nodes" in the "Additional resources" of this section.

25.3.1. Enabling Linux cgroup v1 during installation

You can enable Linux control group version 1 (cgroup v1) when you install a cluster by creating installation manifests.



IMPORTANT

cgroup v1 is a deprecated feature. Deprecated functionality is still included in OpenShift Container Platform and continues to be supported; however, it will be removed in a future release of this product and is not recommended for new deployments.

For the most recent list of major functionality that has been deprecated or removed within OpenShift Container Platform, refer to the *Deprecated and removed features* section of the OpenShift Container Platform release notes.

Procedure

1. Create or edit the **node.config** object to specify the **v1** cgroup:

```
apiVersion: config.openshift.io/v1
kind: Node
metadata:
  name: cluster
spec:
  cgroupMode: "v2"
```

2. Proceed with the installation as usual.

Additional resources

- [OpenShift Container Platform installation overview](#)
- [Configuring the Linux cgroup on your nodes](#)

CHAPTER 26. VALIDATING AN INSTALLATION

You can check the status of an OpenShift Container Platform cluster after an installation by following the procedures in this document.

26.1. REVIEWING THE INSTALLATION LOG

You can review a summary of an installation in the OpenShift Container Platform installation log. If an installation succeeds, the information required to access the cluster is included in the log.

Prerequisites

- You have access to the installation host.

Procedure

- Review the `.openshift_install.log` log file in the installation directory on your installation host:

```
$ cat <install_dir>/.openshift_install.log
```

Example output

Cluster credentials are included at the end of the log if the installation is successful, as outlined in the following example:

```
...
time="2020-12-03T09:50:47Z" level=info msg="Install complete!"
time="2020-12-03T09:50:47Z" level=info msg="To access the cluster as the system:admin
user when using 'oc', run 'export KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'"
time="2020-12-03T09:50:47Z" level=info msg="Access the OpenShift web-console here:
https://console-openshift-console.apps.mycluster.example.com"
time="2020-12-03T09:50:47Z" level=info msg="Login to the console with user: \"kubeadmin\",
and password: \"password\""
time="2020-12-03T09:50:47Z" level=debug msg="Time elapsed per stage:"
time="2020-12-03T09:50:47Z" level=debug msg="  Infrastructure: 6m45s"
time="2020-12-03T09:50:47Z" level=debug msg="Bootstrap Complete: 11m30s"
time="2020-12-03T09:50:47Z" level=debug msg=" Bootstrap Destroy: 1m5s"
time="2020-12-03T09:50:47Z" level=debug msg=" Cluster Operators: 17m31s"
time="2020-12-03T09:50:47Z" level=info msg="Time elapsed: 37m26s"
```

26.2. VIEWING THE IMAGE PULL SOURCE

For clusters with unrestricted network connectivity, you can view the source of your pulled images by using a command on a node, such as `crictl images`.

However, for disconnected installations, to view the source of pulled images, you must review the CRI-O logs to locate the **Trying to access** log entry, as shown in the following procedure. Other methods to view the image pull source, such as the `crictl images` command, show the non-mirrored image name, even though the image is pulled from the mirrored location.

Prerequisites

- You have access to the cluster as a user with the `cluster-admin` role.

Procedure

- Review the CRI-O logs for a master or worker node:

```
$ oc adm node-logs <node_name> -u cri-o
```

Example output

The **Trying to access** log entry indicates where the image is being pulled from.

```
...
Mar 17 02:52:50 ip-10-0-138-140.ec2.internal cri-o[1366]: time="2021-08-05
10:33:21.594930907Z" level=info msg="Pulling image: quay.io/openshift-release-dev/ocp-
release:4.10.0-ppc64le" id=abcd713b-d0e1-4844-ac1c-474c5b60c07c
name=/runtime.v1alpha2.ImageService/PullImage
Mar 17 02:52:50 ip-10-0-138-140.ec2.internal cri-o[1484]: time="2021-03-17
02:52:50.194341109Z" level=info msg="Trying to access \"li0317gcp1.mirror-
registry.qe.gcp.devcluster.openshift.com:5000/ocp/release@sha256:1926eae7cacb9c00f142ec
98b00628970e974284b6ddaf9a6a086cb9af7a6c31\""
Mar 17 02:52:50 ip-10-0-138-140.ec2.internal cri-o[1484]: time="2021-03-17
02:52:50.226788351Z" level=info msg="Trying to access \"li0317gcp1.mirror-
registry.qe.gcp.devcluster.openshift.com:5000/ocp/release@sha256:1926eae7cacb9c00f142ec
98b00628970e974284b6ddaf9a6a086cb9af7a6c31\""
...
```

The log might show the image pull source twice, as shown in the preceding example.

If your **ImageContentSourcePolicy** object lists multiple mirrors, OpenShift Container Platform attempts to pull the images in the order listed in the configuration, for example:

```
Trying to access \"li0317gcp1.mirror-
registry.qe.gcp.devcluster.openshift.com:5000/ocp/release@sha256:1926eae7cacb9c00f142ec
98b00628970e974284b6ddaf9a6a086cb9af7a6c31\"
Trying to access \"li0317gcp2.mirror-
registry.qe.gcp.devcluster.openshift.com:5000/ocp/release@sha256:1926eae7cacb9c00f142ec
98b00628970e974284b6ddaf9a6a086cb9af7a6c31\"
```

26.3. GETTING CLUSTER VERSION, STATUS, AND UPDATE DETAILS

You can view the cluster version and status by running the **oc get clusterversion** command. If the status shows that the installation is still progressing, you can review the status of the Operators for more information.

You can also list the current update channel and review the available cluster updates.

Prerequisites

- You have access to the cluster as a user with the **cluster-admin** role.
- You have installed the OpenShift CLI (**oc**).

Procedure

- Obtain the cluster version and overall status:

-

```
$ oc get clusterversion
```

Example output

```
NAME     VERSION  AVAILABLE  PROGRESSING  SINCE  STATUS
version  4.6.4    True       False        6m25s Cluster version is 4.6.4
```

The example output indicates that the cluster has been installed successfully.

2. If the cluster status indicates that the installation is still progressing, you can obtain more detailed progress information by checking the status of the Operators:

```
$ oc get clusteroperators.config.openshift.io
```

3. View a detailed summary of cluster specifications, update availability, and update history:

```
$ oc describe clusterversion
```

4. List the current update channel:

```
$ oc get clusterversion -o jsonpath='{.items[0].spec}{"\n"}'
```

Example output

```
{"channel":"stable-4.6","clusterID":"245539c1-72a3-41aa-9cec-72ed8cf25c5c"}
```

5. Review the available cluster updates:

```
$ oc adm upgrade
```

Example output

```
Cluster version is 4.6.4
```

```
Updates:
```

```
VERSION IMAGE
```

```
4.6.6 quay.io/openshift-release-dev/ocp-
release@sha256:c7e8f18e8116356701bd23ae3a23fb9892dd5ea66c8300662ef30563d7104f3
9
```

Additional resources

- See [Querying Operator status after installation](#) for more information about querying Operator status if your installation is still progressing.
- See [Troubleshooting Operator issues](#) for information about investigating issues with Operators.
- See [Updating a cluster using the web console](#) for more information on updating your cluster.
- See [Understanding update channels and releases](#) for an overview about update release channels.

26.4. VERIFYING THAT A CLUSTER USES SHORT-TERM CREDENTIALS

You can verify that a cluster uses short-term security credentials for individual components by checking the Cloud Credential Operator (CCO) configuration and other values in the cluster.

Prerequisites

- You deployed an OpenShift Container Platform cluster using the Cloud Credential Operator utility (**ccoctl**) to implement short-term credentials.
- You installed the OpenShift CLI (**oc**).
- You are logged in as a user with **cluster-admin** privileges.

Procedure

- Verify that the CCO is configured to operate in manual mode by running the following command:

```
$ oc get cloudcredentials cluster \
  -o=jsonpath={.spec.credentialsMode}
```

The following output confirms that the CCO is operating in manual mode:

Example output

```
Manual
```

- Verify that the cluster does not have **root** credentials by running the following command:

```
$ oc get secrets \
  -n kube-system <secret_name>
```

where **<secret_name>** is the name of the root secret for your cloud provider.

Platform	Secret name
Amazon Web Services (AWS)	aws-creds
Microsoft Azure	azure-credentials
Google Cloud Platform (GCP)	gcp-credentials

An error confirms that the root secret is not present on the cluster.

Example output for an AWS cluster

```
Error from server (NotFound): secrets "aws-creds" not found
```

- Verify that the components are using short-term security credentials for individual components by running the following command:

```
-
```



```
$ oc get authentication cluster \
  -o jsonpath \
  --template='{ .spec.serviceAccountIssuer }'
```

This command displays the value of the **.spec.serviceAccountIssuer** parameter in the cluster **Authentication** object. An output of a URL that is associated with your cloud provider indicates that the cluster is using manual mode with short-term credentials that are created and managed from outside of the cluster.

- Azure clusters: Verify that the components are assuming the Azure client ID that is specified in the secret manifests by running the following command:

```
$ oc get secrets \
  -n openshift-image-registry installer-cloud-credentials \
  -o jsonpath='{.data}'
```

An output that contains the **azure_client_id** and **azure_federated_token_file** fields confirms that the components are assuming the Azure client ID.

- Azure clusters: Verify that the pod identity webhook is running by running the following command:

```
$ oc get pods \
  -n openshift-cloud-credential-operator
```

Example output

NAME	READY	STATUS	RESTARTS	AGE
cloud-credential-operator-59cf744f78-r8pbq	2/2	Running	2	71m
pod-identity-webhook-548f977b4c-859lz	1/1	Running	1	70m

26.5. QUERYING THE STATUS OF THE CLUSTER NODES BY USING THE CLI

You can verify the status of the cluster nodes after an installation.

Prerequisites

- You have access to the cluster as a user with the **cluster-admin** role.
- You have installed the OpenShift CLI (**oc**).

Procedure

1. List the status of the cluster nodes. Verify that the output lists all of the expected control plane and compute nodes and that each node has a **Ready** status:

```
$ oc get nodes
```

Example output

NAME	STATUS	ROLES	AGE	VERSION
------	--------	-------	-----	---------

```

compute-1.example.com    Ready  worker  33m  v1.29.4
control-plane-1.example.com Ready  master  41m  v1.29.4
control-plane-2.example.com Ready  master  45m  v1.29.4
compute-2.example.com    Ready  worker  38m  v1.29.4
compute-3.example.com    Ready  worker  33m  v1.29.4
control-plane-3.example.com Ready  master  41m  v1.29.4

```

2. Review CPU and memory resource availability for each cluster node:

```
$ oc adm top nodes
```

Example output

```

NAME                CPU(cores) CPU%  MEMORY(bytes) MEMORY%
compute-1.example.com 128m      8%    1132Mi      16%
control-plane-1.example.com 801m     22%   3471Mi      23%
control-plane-2.example.com 1718m    49%   6085Mi      40%
compute-2.example.com  935m     62%   5178Mi      75%
compute-3.example.com  111m     7%    1131Mi      16%
control-plane-3.example.com 942m     26%   4100Mi      27%

```

Additional resources

- See [Verifying node health](#) for more details about reviewing node health and investigating node issues.

26.6. REVIEWING THE CLUSTER STATUS FROM THE OPENSIFT CONTAINER PLATFORM WEB CONSOLE

You can review the following information in the **Overview** page in the OpenShift Container Platform web console:

- The general status of your cluster
- The status of the control plane, cluster Operators, and storage
- CPU, memory, file system, network transfer, and pod availability
- The API address of the cluster, the cluster ID, and the name of the provider
- Cluster version information
- Cluster update status, including details of the current update channel and available updates
- A cluster inventory detailing node, pod, storage class, and persistent volume claim (PVC) information
- A list of ongoing cluster activities and recent events

Prerequisites

- You have access to the cluster as a user with the **cluster-admin** role.

Procedure

- In the **Administrator** perspective, navigate to **Home** → **Overview**.

26.7. REVIEWING THE CLUSTER STATUS FROM RED HAT OPENSIFT CLUSTER MANAGER

From the OpenShift Container Platform web console, you can review detailed information about the status of your cluster on OpenShift Cluster Manager.

Prerequisites

- You are logged in to [OpenShift Cluster Manager](#).
- You have access to the cluster as a user with the **cluster-admin** role.

Procedure

1. Go to the **Clusters** list in [OpenShift Cluster Manager](#) and locate your OpenShift Container Platform cluster.
2. Click the **Overview** tab for your cluster.
3. Review the following information about your cluster:
 - vCPU and memory availability and resource usage
 - The cluster ID, status, type, region, and the provider name
 - Node counts by node type
 - Cluster version details, the creation date of the cluster, and the name of the cluster owner
 - The life cycle support status of the cluster
 - Subscription information, including the service level agreement (SLA) status, the subscription unit type, the production status of the cluster, the subscription obligation, and the service level

TIP

To view the history for your cluster, click the **Cluster history** tab.

4. Navigate to the **Monitoring** page to review the following information:
 - A list of any issues that have been detected
 - A list of alerts that are firing
 - The cluster Operator status and version
 - The cluster's resource usage
5. Optional: You can view information about your cluster that Red Hat Insights collects by navigating to the **Overview** menu. From this menu you can view the following information:
 - Potential issues that your cluster might be exposed to, categorized by risk level

- Health-check status by category

Additional resources

- See [Using Insights to identify issues with your cluster](#) for more information about reviewing potential issues with your cluster.

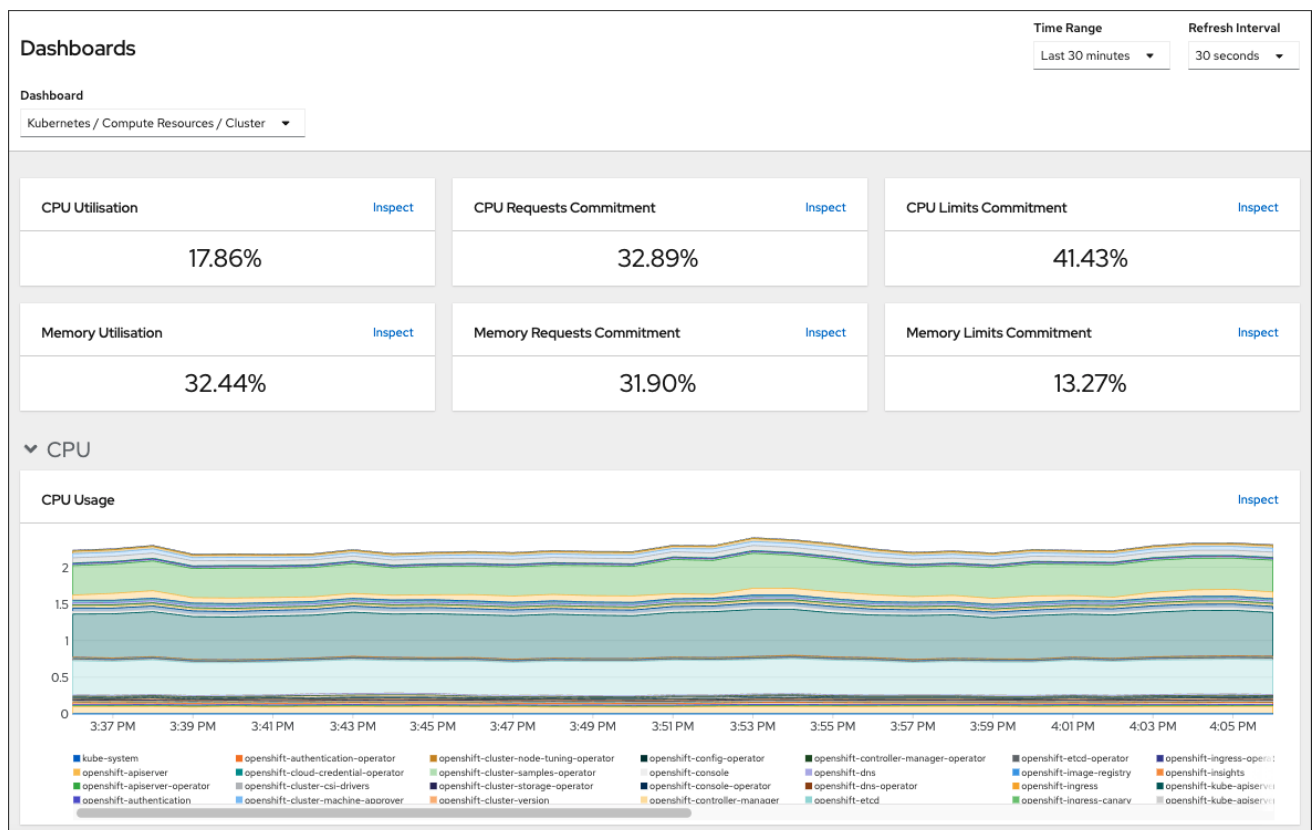
26.8. CHECKING CLUSTER RESOURCE AVAILABILITY AND UTILIZATION

OpenShift Container Platform provides a comprehensive set of monitoring dashboards that help you understand the state of cluster components.

In the **Administrator** perspective, you can access dashboards for core OpenShift Container Platform components, including:

- etcd
- Kubernetes compute resources
- Kubernetes network resources
- Prometheus
- Dashboards relating to cluster and node performance

Figure 26.1. Example compute resources dashboard



Prerequisites

- You have access to the cluster as a user with the **cluster-admin** role.

Procedure

1. In the **Administrator** perspective in the OpenShift Container Platform web console, navigate to **Observe → Dashboards**.
2. Choose a dashboard in the **Dashboard** list. Some dashboards, such as the **etcd** dashboard, produce additional sub-menus when selected.
3. Optional: Select a time range for the graphs in the **Time Range** list.
 - Select a pre-defined time period.
 - Set a custom time range by selecting **Custom time range** in the **Time Range** list.
 - a. Input or select the **From** and **To** dates and times.
 - b. Click **Save** to save the custom time range.
4. Optional: Select a **Refresh Interval**
5. Hover over each of the graphs within a dashboard to display detailed information about specific items.

Additional resources

- See [Monitoring overview](#) for more information about the OpenShift Container Platform monitoring stack.

26.9. LISTING ALERTS THAT ARE FIRING

Alerts provide notifications when a set of defined conditions are true in an OpenShift Container Platform cluster. You can review the alerts that are firing in your cluster by using the Alerting UI in the OpenShift Container Platform web console.

Prerequisites

- You have access to the cluster as a user with the **cluster-admin** role.

Procedure

1. In the **Administrator** perspective, navigate to the **Observe → Alerting → Alerts** page.
2. Review the alerts that are firing, including their **Severity**, **State**, and **Source**.
3. Select an alert to view more detailed information in the **Alert Details** page.

Additional resources

- See [Managing alerts](#) for further details about alerting in OpenShift Container Platform.

26.10. NEXT STEPS

- See [Troubleshooting installations](#) if you experience issues when installing your cluster.

- After installing OpenShift Container Platform, you can [further expand and customize your cluster](#).

CHAPTER 27. TROUBLESHOOTING INSTALLATION ISSUES

To assist in troubleshooting a failed OpenShift Container Platform installation, you can gather logs from the bootstrap and control plane machines. You can also get debug information from the installation program. If you are unable to resolve the issue using the logs and debug information, see [Determining where installation issues occur](#) for component-specific troubleshooting.



NOTE

If your OpenShift Container Platform installation fails and the debug output or logs contain network timeouts or other connectivity errors, review the guidelines for [configuring your firewall](#). Gathering logs from your firewall and load balancer can help you diagnose network-related errors.

27.1. PREREQUISITES

- You attempted to install an OpenShift Container Platform cluster and the installation failed.

27.2. GATHERING LOGS FROM A FAILED INSTALLATION

If you gave an SSH key to your installation program, you can gather data about your failed installation.



NOTE

You use a different command to gather logs about an unsuccessful installation than to gather logs from a running cluster. If you must gather logs from a running cluster, use the **oc adm must-gather** command.

Prerequisites

- Your OpenShift Container Platform installation failed before the bootstrap process finished. The bootstrap node is running and accessible through SSH.
- The **ssh-agent** process is active on your computer, and you provided the same SSH key to both the **ssh-agent** process and the installation program.
- If you tried to install a cluster on infrastructure that you provisioned, you must have the fully qualified domain names of the bootstrap and control plane nodes.

Procedure

1. Generate the commands that are required to obtain the installation logs from the bootstrap and control plane machines:
 - If you used installer-provisioned infrastructure, change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install gather bootstrap --dir <installation_directory> 1
```

- 1** **installation_directory** is the directory you specified when you ran **./openshift-install create cluster**. This directory contains the OpenShift Container Platform definition files that the installation program creates.

For installer-provisioned infrastructure, the installation program stores information about the cluster, so you do not specify the hostnames or IP addresses.

- If you used infrastructure that you provisioned yourself, change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install gather bootstrap --dir <installation_directory> \ 1
--bootstrap <bootstrap_address> \ 2
--master <master_1_address> \ 3
--master <master_2_address> \ 4
--master <master_3_address>" 5
```

- 1** For **installation_directory**, specify the same directory you specified when you ran **./openshift-install create cluster**. This directory contains the OpenShift Container Platform definition files that the installation program creates.
- 2** **<bootstrap_address>** is the fully qualified domain name or IP address of the cluster's bootstrap machine.
- 3** **4** **5** For each control plane, or master, machine in your cluster, replace **<master*_address>** with its fully qualified domain name or IP address.



NOTE

A default cluster contains three control plane machines. List all of your control plane machines as shown, no matter how many your cluster uses.

Example output

```
INFO Pulling debug logs from the bootstrap machine
INFO Bootstrap gather logs captured here "<installation_directory>/log-bundle-
<timestamp>.tar.gz"
```

If you open a Red Hat support case about your installation failure, include the compressed logs in the case.

27.3. MANUALLY GATHERING LOGS WITH SSH ACCESS TO YOUR HOST(S)

Manually gather logs in situations where **must-gather** or automated collection methods do not work.



IMPORTANT

By default, SSH access to the OpenShift Container Platform nodes is disabled on the Red Hat OpenStack Platform (RHOSP) based installations.

Prerequisites

- You must have SSH access to your host(s).

Procedure

1. Collect the **bootkube.service** service logs from the bootstrap host using the **journalctl** command by running:

```
$ journalctl -b -f -u bootkube.service
```

2. Collect the bootstrap host's container logs using the podman logs. This is shown as a loop to get all of the container logs from the host:

```
$ for pod in $(sudo podman ps -a -q); do sudo podman logs $pod; done
```

3. Alternatively, collect the host's container logs using the **tail** command by running:

```
# tail -f /var/lib/containers/storage/overlay-containers/*/userdata/ctr.log
```

4. Collect the **kubelet.service** and **crio.service** service logs from the master and worker hosts using the **journalctl** command by running:

```
$ journalctl -b -f -u kubelet.service -u crio.service
```

5. Collect the master and worker host container logs using the **tail** command by running:

```
$ sudo tail -f /var/log/containers/*
```

27.4. MANUALLY GATHERING LOGS WITHOUT SSH ACCESS TO YOUR HOST(S)

Manually gather logs in situations where **must-gather** or automated collection methods do not work.

If you do not have SSH access to your node, you can access the systems journal to investigate what is happening on your host.

Prerequisites

- Your OpenShift Container Platform installation must be complete.
- Your API service is still functional.
- You have system administrator privileges.

Procedure

1. Access **journal** unit logs under **/var/log** by running:

```
$ oc adm node-logs --role=master -u kubelet
```

2. Access host file paths under **/var/log** by running:

```
$ oc adm node-logs --role=master --path=openshift-apiserver
```

27.5. GETTING DEBUG INFORMATION FROM THE INSTALLATION PROGRAM

You can use any of the following actions to get debug information from the installation program.

- Look at debug messages from a past installation in the hidden **.openshift_install.log** file. For example, enter:

```
$ cat ~/<installation_directory>/.openshift_install.log 1
```

- 1 For **installation_directory**, specify the same directory you specified when you ran **./openshift-install create cluster**.

- Change to the directory that contains the installation program and re-run it with **--log-level=debug**:

```
$ ./openshift-install create cluster --dir <installation_directory> --log-level debug 1
```

- 1 For **installation_directory**, specify the same directory you specified when you ran **./openshift-install create cluster**.

27.6. REINSTALLING THE OPENSIFT CONTAINER PLATFORM CLUSTER

If you are unable to debug and resolve issues in the failed OpenShift Container Platform installation, consider installing a new OpenShift Container Platform cluster. Before starting the installation process again, you must complete thorough cleanup. For a user-provisioned infrastructure (UPI) installation, you must manually destroy the cluster and delete all associated resources. The following procedure is for an installer-provisioned infrastructure (IPI) installation.

Procedure

1. Destroy the cluster and remove all the resources associated with the cluster, including the hidden installer state files in the installation directory:

```
$ ./openshift-install destroy cluster --dir <installation_directory> 1
```

- 1 **installation_directory** is the directory you specified when you ran **./openshift-install create cluster**. This directory contains the OpenShift Container Platform definition files that the installation program creates.

2. Before reinstalling the cluster, delete the installation directory:

```
$ rm -rf <installation_directory>
```

3. Follow the procedure for installing a new OpenShift Container Platform cluster.

Additional resources

- [Installing an OpenShift Container Platform cluster](#)

CHAPTER 28. SUPPORT FOR FIPS CRYPTOGRAPHY

You can install an OpenShift Container Platform cluster in FIPS mode.

OpenShift Container Platform is designed for FIPS. When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the x86_64, ppc64le, and s390x architectures.

For more information about the NIST validation program, see [Cryptographic Module Validation Program](#). For the latest NIST status for the individual versions of RHEL cryptographic libraries that have been submitted for validation, see [Compliance Activities and Government Standards](#).



IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a RHEL 9 computer that is configured to operate in FIPS mode, and you must use a FIPS-capable version of the installation program. See the section titled *Obtaining a FIPS-capable installation program using `oc adm extract`*.

For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#).

For the Red Hat Enterprise Linux CoreOS (RHCOS) machines in your cluster, this change is applied when the machines are deployed based on the status of an option in the **install-config.yaml** file, which governs the cluster options that a user can change during cluster deployment. With Red Hat Enterprise Linux (RHEL) machines, you must enable FIPS mode when you install the operating system on the machines that you plan to use as worker machines.

Because FIPS must be enabled before the operating system that your cluster uses boots for the first time, you cannot enable FIPS after you deploy a cluster.

28.1. OBTAINING A FIPS-CAPABLE INSTALLATION PROGRAM USING OC ADM EXTRACT

OpenShift Container Platform requires the use of a FIPS-capable installation binary to install a cluster in FIPS mode. You can obtain this binary by extracting it from the release image by using the OpenShift CLI (**oc**). After you have obtained the binary, you proceed with the cluster installation, replacing all instances of the **openshift-install** command with **openshift-install-fips**.

Prerequisites

- You have installed the OpenShift CLI (**oc**) with version 4.16 or newer.

Procedure

- Extract the FIPS-capable binary from the installation program by running the following command:

```
$ oc adm release extract --registry-config "${pullsecret_file}" --command=openshift-install-fips --to "${extract_dir}" ${RELEASE_IMAGE}
```

where:

<pullsecret_file>

Specifies the name of a file that contains your pull secret.

<extract_dir>

Specifies the directory where you want to extract the binary.

<RELEASE_IMAGE>

Specifies the Quay.io URL of the OpenShift Container Platform release you are using. For more information on finding the release image, see *Extracting the OpenShift Container Platform installation program*.

2. Proceed with cluster installation, replacing all instances of the **openshift-install** command with **openshift-install-fips**.

Additional resources

- [Extracting the OpenShift Container Platform installation program](#)

28.2. OBTAINING A FIPS-CAPABLE INSTALLATION PROGRAM USING THE PUBLIC OPENSHIFT MIRROR

OpenShift Container Platform requires the use of a FIPS-capable installation binary to install a cluster in FIPS mode. You can obtain this binary by downloading it from the public OpenShift mirror. After you have obtained the binary, proceed with the cluster installation, replacing all instances of the **openshift-install** binary with **openshift-install-fips**.

Prerequisites

- You have access to the internet.

Procedure

1. Download the installation program from <https://mirror.openshift.com/pub/openshift-v4/clients/ocp/latest-4.16/openshift-install-rhel9-amd64.tar.gz>.
2. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-rhel9-amd64.tar.gz
```

3. Proceed with cluster installation, replacing all instances of the **openshift-install** command with **openshift-install-fips**.

28.3. FIPS VALIDATION IN OPENSHIFT CONTAINER PLATFORM

OpenShift Container Platform uses certain FIPS validated or Modules In Process modules within RHEL and RHCOS for the operating system components that it uses. See [RHEL core crypto components](#). For example, when users use SSH to connect to OpenShift Container Platform clusters and containers, those connections are properly encrypted.

OpenShift Container Platform components are written in Go and built with Red Hat's golang compiler. When you enable FIPS mode for your cluster, all OpenShift Container Platform components that require cryptographic signing call RHEL and RHCOS cryptographic libraries.

Table 28.1. FIPS mode attributes and limitations in OpenShift Container Platform 4.16

Attributes	Limitations
FIPS support in RHEL 9 and RHCOS operating systems.	The FIPS implementation does not use a function that performs hash computation and signature generation or validation in a single step. This limitation will continue to be evaluated and improved in future OpenShift Container Platform releases.
FIPS support in CRI-O runtimes.	
FIPS support in OpenShift Container Platform services.	
FIPS validated or Modules In Process cryptographic module and algorithms that are obtained from RHEL 9 and RHCOS binaries and images.	
Use of FIPS compatible golang compiler.	TLS FIPS support is not complete but is planned for future OpenShift Container Platform releases.
FIPS support across multiple architectures.	FIPS is currently only supported on OpenShift Container Platform deployments using x86_64 , ppc64le , and s390x architectures.

28.4. FIPS SUPPORT IN COMPONENTS THAT THE CLUSTER USES

Although the OpenShift Container Platform cluster itself uses FIPS validated or Modules In Process modules, ensure that the systems that support your OpenShift Container Platform cluster use FIPS validated or Modules In Process modules for cryptography.

28.4.1. etcd

To ensure that the secrets that are stored in etcd use FIPS validated or Modules In Process encryption, boot the node in FIPS mode. After you install the cluster in FIPS mode, you can [encrypt the etcd data](#) by using the FIPS-approved **aes cbc** cryptographic algorithm.

28.4.2. Storage

For local storage, use RHEL-provided disk encryption or Container Native Storage that uses RHEL-provided disk encryption. By storing all data in volumes that use RHEL-provided disk encryption and enabling FIPS mode for your cluster, both data at rest and data in motion, or network data, are protected by FIPS validated or Modules In Process encryption. You can configure your cluster to encrypt the root filesystem of each node, as described in [Customizing nodes](#).

28.4.3. Runtimes

To ensure that containers know that they are running on a host that is using FIPS validated or Modules In Process cryptography modules, use CRI-O to manage your runtimes.

28.5. INSTALLING A CLUSTER IN FIPS MODE


To install a cluster in FIPS mode, follow the instructions to install a customized cluster on your preferred infrastructure. Ensure that you set **fips: true** in the **install-config.yaml** file before you deploy your cluster.



IMPORTANT

To enable FIPS mode for your cluster, you must run the installation program from a RHEL computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see [Installing the system in FIPS mode](#) .

- [Amazon Web Services](#)
- [Microsoft Azure](#)
- [Bare metal](#)
- [Google Cloud Platform](#)
- [IBM Cloud®](#)
- [IBM Power®](#)
- [IBM Z® and IBM® LinuxONE](#)
- [IBM Z® and IBM® LinuxONE with RHEL KVM](#)
- [Red Hat OpenStack Platform \(RHOSP\)](#)
- [VMware vSphere](#)



NOTE

If you are using Azure File storage, you cannot enable FIPS mode.

To apply **AES CBC** encryption to your etcd data store, follow the [Encrypting etcd data](#) process after you install your cluster.

If you add RHEL nodes to your cluster, ensure that you enable FIPS mode on the machines before their initial boot. See [Adding RHEL compute machines to an OpenShift Container Platform cluster](#) and [Installing the system in FIPS mode](#) .